Tight Bounds for MIS in Multichannel Radio Networks

Sebastian Daum
Dept. of Comp. Science
U. of Freiburg, Germany
daum.sebastian@gmail.com

Fabian Kuhn
Dept. of Comp. Science
U. of Freiburg, Germany
kuhn@cs.uni-freiburg.de

Abstract

In [8] an algorithm has been presented that computes a maximal independent set (MIS) within $O(\log^2 n/\mathcal{F} + \log n \operatorname{polyloglog} n)$ rounds in an n-node multichannel radio network with $\mathcal{F}$ communication channels. The paper uses a multichannel variant of the standard graph-based radio network model without collision detection and it assumes that the network graph is a polynomially bounded independence graph (BIG), a natural combinatorial generalization of well-known geographic families. The upper bound of [8] is known to be optimal up to a $\operatorname{polyloglog}$ factor.

In this paper, we adapt algorithm and analysis to improve the result of [8] in two ways. Mainly, we get rid of the $\operatorname{polyloglog}$ factor in the runtime and we thus obtain an asymptotically optimal multichannel radio network MIS algorithm. In addition, our new analysis allows to generalize the class of graphs from those with polynomially bounded local independence to graphs where the local independence is bounded by an arbitrary function of the neighborhood radius.

Keywords: maximal independent set, radio network, multichannel, shared spectrum, growth-bounded graph, bounded independence graph

1 Introduction

In recent years there has been an increased interest in algorithms for *shared spectrum networks* [26]. Nowadays, most modern wireless communication networks feature a multitude of communication frequencies [1, 2, 5][1]—and we can certainly expect this trend to continue.

In the light of this development, in the present paper, we settle the question of determining the *optimal* asymptotic time complexity of computing a maximal independent set (MIS) in the multichannel variant of the classic radio network model first introduced in [4, 7]. The task of constructing an MIS is one of the best studied problems in the area of large-scale wireless networks. On the one hand this is due to the fact that MIS (together with coloring problems) is one of the key problems to study the problem of symmetry breaking in large, decentralized systems. On the other hand an MIS provides a simple local clustering of the graph, which can be used as a building block for computing more enhanced organization structures in these networks such as, e.g., a communication backbone based on a connected dominating set [6, 18, 29]. This is specifically relevant in the context of wireless mobile ad hoc networks or sensor networks, in which devices cannot rely on already existing infrastructure to organize themselves—devices need to compute a meaningful structure by themselves to coordinate their interactions.

Related Work. In [3, 21] Alon et al. and Luby presented a simple and efficient randomized parallel algorithm to compute and MIS of a general graph. It is straightforward to a standard distributed message passing model and as a consequence, the algorithm soon became an archetype for many distributed MIS algorithms also in other—usually more limiting—settings. The model we assume here is an extension to the radio network model, for which an MIS algorithm with runtime $O(\log^2 n)$ has been presented in [22] for the class of unit disk graphs (UDGs). This algorithm has been proven to be asymptotically optimal [11] even for more basic version of the problem known as the wake-up problem in single-hop radio networks. While the UDG restriction is well-known and popular, a more general variant known as growth-bounded graphs or bounded independence graphs that contains UDGs has also become the focus of quite some research, e.g., [20, 24, 25]. In particular, in [25] it is shown that an MIS and many related structures can be compute in (asymptotically optimal) $O(\log^* n)$ rounds in such graphs.

Much of the early algorithmic research on multichannel radio networks has focused on networks with faults assuming a malicious adversary that can jam up to t of the $\mathcal{F}$ available channels [10, 13–17, 19, 27, 28]. In addition, for fault-free networks, in [23] a series of lower bound proofs have been provided, which show that $\Omega(\log^2 n/\mathcal{F} + \log n)$ rounds are needed to solve any problem which requires communication. In [10] a new technique (called heralding) to deal with congestion in multichannel radio networks has been established to solve leader election in single-hop networks in time asymptotically matching the lower bound of [23]. This technique has been extended in [12] and [8] to solve the problems of computing an approximate minimum dominating set and an MIS, respectively. Our research here is based on this work and in particular on the MIS algorithm of [8].

Contributions. In radio network models, in almost all cases a restriction to the underlying graph model is being assumed. One of the most general ones are so-called α-bounded independence graphs, where $\alpha(r)$ is a function that limits the size of a *maximum* independent set in any r-neighborhood of the given graph. The MIS algorithm from [8] solves the MIS problem in time $O(\log^2 n/\mathcal{F} + \log n(\log\log n)^d)$ in such graphs for which α is bounded by polynomial of degree d. Here we get rid of the $\mathrm{polyloglog}$ factor and thus show how to close the gap to the lower bound from [23]. At the same time, we remove any restriction on the function α. We do so by adjusting the algorithm from [8]—and though the change in the algorithm is relatively small, it leads to a significantly more involved analysis.

[1]For example, the IEEE 802.11 WLAN standard provides a channel spectrum of up to 200 (partially overlapping) channels and Bluetooth specifies 79 usable channels.

2 Preliminaries

This paper bases strongly on [8] and [9], the former being the proceedings version and the latter the complete version. However, we try to be as self-contained as possible.

Radio Network Model. We model the network as an n-node graph $G = (V, E)$. We assume that n or a polynomial upper bound on n is known by all nodes. Nodes start out dormant and are awakened/activated by an adversary. While nodes do not have access to a global clock, communication is assumed to happen in synchronous time slots (rounds). The network comprises $\mathcal{F}$ communication channels. In each round each node can choose to operate on one channel, either by listening or broadcasting. A node that broadcasts does not receive any message in that round, and its signal reaches all neighbors that operate on the same channel. A node v listeing on some channel can decode an incoming message iff in the given round, exactly one of its neighbors broadcasts on the same channel. If two or more neighbors broadcast, their signals collide at v and v receives nothing, unable to detect this collision. A node can only operate on one channel in each round and therefore it does not learn anything about events on other channels.

Notation. In our algorithm all nodes move between a finite set of states: $\mathbb{W}$ – *waiting*, $\mathbb{D}$ – *decay*, $\mathbb{A}$ – *active*, $\mathbb{H}'$ – *herald candidate*, $\mathbb{H}$ – *herald*, $\mathbb{L}'$ – *leader candidate*, $\mathbb{L}$ – *leader*, $\mathbb{M}$ – *MIS node*, $\mathbb{E}$ – *eliminated/dominated*. We overload this notation to also indicate the set of nodes being in a certain state, e.g., $\mathbb{A} := \{v \in V : v \text{ is in state } \mathbb{A}\}$. Since nodes change their states, in case of ambiguity, we write $\mathbb{A}_r$ for the set of active nodes in round r. State changes always happen between rounds. We define $V_{\text{hf}} := \mathbb{A} \cup \mathbb{H}' \cup \mathbb{L}' \cup \mathbb{H} \cup \mathbb{L}$ as the nodes in the so-called herald filter.

We use $N(v)$ to denote the neighbors of v in G, while we use $N^k(v)$ to denote the set of nodes in distance at most k from v, including v itself. We also often write $N_{\mathbb{S}}(u)$ or $N^k_{\mathbb{S}}(u)$ to abbreviate $N(u) \cap \mathbb{S}$ or $N^k(u) \cap \mathbb{S}$ respectively, for some state $\mathbb{S}$. For $S \subseteq V$ we let $N(S) := \bigcup_{v \in S} N(v)$. We call a node v *alone* or *lonely*, if $N_{V_{\text{hf}} \cup \mathbb{M}}(u) = \emptyset$.

We say that an event A happens *with high probability (w.h.p.)*, *with decent probability*, or *with constant probability (w.c.p.)*, if it happens with probability at least $1 - n^{-c}$, $1 - \log^{-c} n$, or $\Omega(1)$, respectively, where c is a constant that can be chosen arbitrarily large. By $x \gg y$ we denote that $x > cy$ for a sufficiently large $c > 1$.

Bounded Independence. In addition to the communication characteristics of the network, we require the network graph to be a bounded independence graph (BIG) [20, 24]. A graph G is called an *α-bounded independence graph* with *independence function* $\alpha : \mathbb{N} \to \mathbb{N}$, if for every node v. no independent set S of the subgraph of G induced by $N^d(v)$ exceeds cardinality $\alpha(d)$. Note that in particular, α does not depend on n and thus for every fixed d, $\alpha(d)$ is a constant. In [8], α is required to be a polynomial, whereas in this paper, we put no restrictions on α. It can easily be verified that one can always upper bound the largest independent set of the subgraph induced by $N^d(v)$ by $\alpha(2)^d$ and thus any independence function is always upper bounded by some exponential function. For simplicity we define a constant $\alpha := \alpha(2)$ and we assume that all nodes know the value of α.

Number of Channels. We assume that $\mathcal{F} = \omega(1)$ as otherwise single channel algorithms achieve the same asymptotic time bounds. For $\mathcal{F} = \omega(\log n)$ we only actually use $\Theta(\log n)$ channels since more channels do not lead to an additional asymptotic advantage. For ease of exposition we assume $\mathcal{F} = \Omega(\log \log n)$ and refer to [9] for an explanation of how to adapt the algorithm for the case $\mathcal{F} = o(\log \log n)$.

Maximal Independent Set. We say an algorithm computes an MIS in time T, if the following properties hold w.h.p. for each round r and node v (waking up in round r_v):

(P1) v declares itself as either *dominating* ($\in \mathbb{M}$) or *dominated* ($\in \mathbb{E}$) before round $r_v + T$ and this decision is permanent.

(P2) If v is *dominated* in round r, then $N(v) \cap \mathbb{M}_r \neq \emptyset$.

(P3) If v is *dominating* in round r, then $N(v) \cap \mathbb{M}_r = \emptyset$.

3 Algorithm Description

Algorithm 1 HeraldMIS—core structure

Input: $\sigma_{\oplus}$, $\sigma_{\ominus}$, Δ_{max}, π_{ℓ}, α, n, $n_{\mathcal{D}} = \Theta(\mathcal{F})$, $n_{\mathcal{A}} = \Theta(\log\log n)$, $n_{\mathcal{R}} = \Theta(\alpha(2))$,
$\tau_{\mathbb{W}} = \Theta(\log n)$, $\tau_{\mathbb{D}} = \Theta(\log n/\mathcal{F})$, $\tau_{lonely} = \Theta(\log^2 n/\mathcal{F} + \log n)$, $\tau_{red\text{-}blue} = \Theta(\log n)$

States: $\mathbb{W}$—waiting, $\mathbb{D}$—decay, $\mathbb{M}$—MIS node, $\mathbb{E}$—eliminated
$\mathbb{A}$—active, $\mathbb{L}/\mathbb{L}'$—leader (candidate), $\mathbb{H}/\mathbb{H}'$—herald (candidate)

Channels: $\mathcal{R}_1, \ldots, \mathcal{R}_{n_{\mathcal{R}}}$—report, $\mathcal{D}_1, \ldots, \mathcal{D}_{n_{\mathcal{D}}}$—decay,
$\mathcal{A}_1, \ldots, \mathcal{A}_{n_{\mathcal{A}}}$—herald, $\mathcal{H}$—handshake, $\mathcal{G}$—red-blue game

count $\leftarrow 0$; *state* $\leftarrow \mathbb{W}$; $\gamma \leftarrow \perp$; *lonely* $\leftarrow \perp$; $\gamma_{min} \leftarrow \log^{-24} n$
while *state* $\neq \mathbb{E}$ **do**
  count $\leftarrow$ *count* $+ 1$
  lonely $\leftarrow$ *lonely* $+ 1$
  $\gamma \leftarrow \min\{\gamma \cdot \sigma_{\oplus}, 1/2\}$
  uniformly at random pick $q \in [0, 1)$, $j \in \{1, \ldots, n_{\mathcal{D}}\}$ and $k \in \{1, \ldots, n_{\mathcal{R}}\}$
  switch *state* **do**
    case $\mathbb{W}$ or $\mathbb{D}$: run DFilter ▷ stage 1—decay filter
    case $\mathbb{A}$: run HeraldProtocol ▷ stage 2—herald filter
    case $\mathbb{H}'$ or $\mathbb{L}'$: run Handshake
    case $\mathbb{H}$ or $\mathbb{L}$: run RedBlueGame
    case $\mathbb{M}$: run Dominator ▷ stage 3—MIS node
  if *lonely* $= \tau_{lonely}$ **then**
    state $\leftarrow \mathbb{M}$
endWhile

Theorem 3.1. *Algorithm* HeraldMIS *solves MIS within* $O(\log^2 n/\mathcal{F} + \log n)$ *rounds.*

We first give a short summary of how the algorithm works, which includes a recap of results from [9]. The algorithm is divided into three stages, the decay filter (states $\mathbb{W}$ and $\mathbb{D}$), the herald filter (states $\mathbb{A}$, $\mathbb{L}'$, $\mathbb{H}'$, $\mathbb{L}$ and $\mathbb{H}$), and decided nodes (states $\mathbb{M}$ and $\mathbb{E}$). Nodes move forward within those stages—possibly omitting the herald filter—but never backwards. The decay filter is a powerful tool (which we use as a black box) that provides that over the full runtime of the algorithm the degree of the graph induced by nodes in the herald filter is bounded by $O(\log^3 n)$. In short, nodes first only listen for a while ($\mathbb{W}$), then they start broadcasting on one out of $\Theta(\mathcal{F})$ random channels with probability $1/n$ ($\mathbb{D}$), doubling this probability every $O(\log n/\mathcal{F})$ rounds. A node that broadcasts moves to the herald filter and a node that receives a message restarts with $\mathbb{W}$. The decay filter has not changed and for a detailed analysis we refer to [9], while pseudo-code is given in Algorithm 2.

Eliminated nodes ($\mathbb{E}$) know that they have a neighbor in the MIS and stop their protocol. MIS nodes ($\mathbb{M}$) try to inform their neighborhood (eliminating them), but they also actively disrupt protocols in the herald filter, causing them to fail; for more details confer Algorithm 3. Apart from this, there is no influence between nodes being in different stages.

The focus of this paper is almost exclusively on the herald filter. It helps for understanding the complex algorithm to *only* think of the graph that is induced by nodes in the herald filter and to recall that its maximum degree is polylogarithmic in n.

The herald filter is divided into three blocks, *active state/herald protocol* ($\mathbb{A}$), *handshake protocol* ($\mathbb{L}'$ and $\mathbb{H}'$) and *red-blue protocol* ($\mathbb{L}$ and $\mathbb{H}$). The first block has the purpose of nodes trying to make contact with surrounding nodes. If this does indeed happen, both nodes engage in the handshake, which is only successful, if none of the two nodes neighbors any MIS node or a node in the third block. If the handshake succeeds, both nodes start a series of coin flipping games, with the sole purpose of ensuring that no two

Algorithm 2 DFILTER, run at process v

Input: $\mathcal{F}$, $n_{\mathcal{D}} = \Theta(\mathcal{F})$, $n_{\mathcal{R}} \geq 3\alpha^2$, $\tau_{\mathbb{W}} = \Theta(\log n)$, $\tau_{\mathbb{D}} = \Theta(\log n/\mathcal{F})$
States: $\mathbb{W}$—waiting, $\mathbb{D}$—decay
Channels: $\mathcal{R}_1, \ldots, \mathcal{R}_{n_{\mathcal{R}}}$—report, $\mathcal{D}_1, \ldots, \mathcal{D}_{n_{\mathcal{D}}}$—decay

count $\leftarrow 0$, *state* $\leftarrow \mathbb{W}$
while *state* $\neq \mathbb{E}$ **do**
  count $\leftarrow$ *count* $+ 1$
  pick $i \in \{1, \ldots, n_{\mathcal{R}}\}$, $k \in \{1, \ldots, n_{\mathcal{D}}\}$ and $q \in [0, 1)$ uniformly at random
  switch *state* **do**
    case $\mathbb{W}$
      listen on channel $\mathcal{R}_i$
      if *count* $= \tau_{\mathbb{W}}$ **then**
        count $\leftarrow 0$, *state* $\leftarrow \mathbb{D}$, *phase* $\leftarrow 0$
    case $\mathbb{D}$
      switch q **do**
        case $q \in [0, 2^{\textit{phase}}/n)$
          send *msg* $= (\textit{ID}, \textit{state})$ on $\mathcal{D}_k$
          exit decay filter and enter herald filter
        case $q \in [2^{\textit{phase}}/n, 1/2)$
          listen on $\mathcal{D}_k$
        case $q \in [1/2, 1)$
          listen on $\mathcal{R}_i$
      if *count* $= \tau_{\mathbb{D}}$ **then**
        count $\leftarrow 0$, *phase* $\leftarrow \min\{\textit{phase} + 1, \log n - 2\}$

Upon receiving a message *msg* $= (\textit{\textbf{msg.ID}}, \textit{\textbf{msg.state}})$

if *msg.state* $= \mathbb{D}$ **then** ▷ restart decay filter
  count $\leftarrow 0$, *state* $\leftarrow \mathbb{W}$
if *msg.state* $= \mathbb{M}$ **then**
  state $\leftarrow \mathbb{E}$

Algorithm 3 DOMINATOR

if *enforce* **then**
  send $(\textit{state}, \textit{ID})$ on $\mathcal{H}$
  enforce $\leftarrow$ **false**
else
  switch q **do**
    case $q \in \left[0, \frac{1}{2}\right)$
      send $(\textit{state}, \textit{ID})$ on $\mathcal{H}$
      enforce $\leftarrow$ **false**
    case $q \in \left[\frac{1}{2}, \frac{3}{4}\right)$
      send $(\textit{state}, \textit{ID})$ on $\mathcal{G}$
      enforce $\leftarrow$ **true**
    case $q \in \left[\frac{3}{4}, 1\right)$
      send $(\textit{state}, \textit{ID})$ on $\mathcal{R}_k$
      enforce $\leftarrow$ **true**

nodes, that became leaders ($\mathbb{L}$) *simultaneously*, can join the MIS. The blocks that differ from the algorithm description in [9] are the active state and the red-blue protocol, although changes in the latter are made to impact nodes in the active state.

Ahead of all we want to mention that there are two ways for a node v to join the MIS—either by waiting for a long time without hearing from any nearby node, or by successfully communicating with a node u during the active state, *teaming up* with it (as a leader-herald pair) and together passing through the handshake and the red-blue protocol. The farther a pair of nodes advances in these blocks, the closer its leader is to become an MIS node. We now recap and describe the behavior of a node v in the herald filter, i.e., $v \in V_{\text{hf}}$, pointing out when changes to the original algorithm occur.

Loneliness. v maintains a counter *lonely*. Whenever v hears from another node, it resets *lonely* to zero. If *lonely* ever exceeds $\tau_{lonely} = \Theta\left(\log^2 n/\mathcal{F} + \log n\right)$, then v *assumes* that it is *alone/lonely* in the herald filter (i.e., $N_{V_{\text{hf}} \cup \mathbb{M}}(v) = \emptyset$) and joins the MIS—w.h.p., this action is safe, i.e., should v not be alone, then the neighbors of v are far from becoming MIS nodes themselves and v has enough time to eliminate them.[2]

Activity. Also, v maintains an *activity* value $\gamma(v) \in \left[\gamma_{min}, 1/2\right]$, where γ_{min}, the initial value, is in $\Omega(1/\operatorname{polylog} n)$. γ governs the behavior of v in $\mathbb{A}$, but all nodes in V_{hf} maintain this value. Nodes outside V_{hf} have zero activity. $\gamma(v)$ increases by a (small) constant factor $\sigma_{\oplus} > 1$ each round, such that after $\Theta(\log\log n)$ rounds it would reach the maximum value $1/2$. However, whenever v receives from a neighboring *leader* or *herald*, then v reduces $\gamma(v)$ by a (large) constant factor $\sigma_{\ominus} \gg \sigma_{\oplus}$. This is a change to the original algorithm, where γ could only increase. The reason is the following. Leaders are nodes that likely become MIS nodes, and if they do then they eliminate their neighbors anyway. For safety reasons a leader l needs to wait for $\Theta(\log n)$ rounds before it may join the MIS. During that time, if l's neighbors keep high activity values, progress stagnation can occur in a $\delta' = O(\log\log n)$ neighborhood of l, which is why in [9] an $\alpha(\delta') = O(\operatorname{polyloglog} n)$ speed loss had to be accepted. We show here that by reducing activities, this stagnation can be eliminated. At the same time, 'unjustified' reductions only cause 'minor damage' that can easily be mitigated. This change is reflected in line 16 of Algorithm 4.

Algorithm 4 HERALDPROTOCOL

1: pick an $i \in \{1, \dots, n_{\mathcal{A}}, \perp\}$ randomly with distribution $\mathbf{P}(i = \perp) = 2^{-n_{\mathcal{A}}}$ and $\mathbf{P}(i = j) = 2^{-j}$
2: **if** $i = \perp$ **then** $q = 1$
3: **switch** q **do**
4:     **case** $q \in [0, \pi_\ell \gamma)$
5:         listen on $\mathcal{A}_i$
6:         **if** *msg* $\neq \emptyset$ **then**
7:             $ID_{leader} \leftarrow msg.ID$; $state \leftarrow \mathbb{H}'$; $count \leftarrow 0$; $handshake \leftarrow succ$; $lonely \leftarrow 0$
8:     **case** $q \in [\pi_\ell \gamma, \gamma)$
9:         send (ID) on $\mathcal{A}_i$
10:         $state \leftarrow \mathbb{L}'$; $count \leftarrow 0$; $handshake \leftarrow succ$;
11:     **case** $q \in [\gamma, 1]$
12:         listen on $\mathcal{R}_k$
13:         **if** $msg.state = \mathbb{M}$ **then**
14:             $state \leftarrow \mathbb{E}$; $\gamma \leftarrow 0$
15:         **if** $msg.state \in \{\mathbb{L}, \mathbb{H}\}$ **then**
16:             $\gamma \leftarrow \max\{\gamma/\sigma_{\ominus}, \gamma_{min}\}$; $lonely \leftarrow 0$

Herald Protocol. Confer Algorithm 4 for detailed pseudo-code. A node v in the active state ($\mathbb{A}$) participates in the *herald protocol* with probability $\gamma(v) \in [\gamma_{min}, 1/2]$, otherwise it tries to learn of nearby leaders, heralds or MIS nodes, by listening to one of *constant* many *report channels* $\mathcal{R}_1, \dots, \mathcal{R}_{n_{\mathcal{R}}}$, $n_{\mathcal{R}} \geq 3\alpha^2$. If v participates in the herald protocol, then it chooses a channel $\mathcal{A}_i$ from $\mathcal{A}_1, \dots, \mathcal{A}_{n_{\mathcal{A}}}$ with probability 2^{-i}.

[2]In [9] there existed some component called *loneliness support block*, operating on its own set of channels $\mathcal{S}_1, \dots, \mathcal{S}_{n_{\mathcal{S}}}$; this block and its channels have been removed.

It then listens on $\mathcal{A}_i$ with probability $\pi_\ell \leq 1/10$ or broadcasts its *ID* otherwise.[3] If v listens, but receives nothing, nothing happens and v stays in $\mathbb{A}$. Should v receive the message of another node u on $\mathcal{A}_i$, then next round it engages with u in the *handshake protocol* as a *herald candidate* ($\mathbb{H}'$), in the hope of moving forward to the *red-blue protocol* together with u. Should v choose to broadcast, then it deterministically pursues the handshake as a *leader candidate* ($\mathbb{L}'$), hoping that some other node u has heard its message and joins in for the handshake.

Algorithm 5 HANDSHAKE

1: **switch** *state* **do**
2:   **case** $\mathbb{H}'$
3:     **switch** *count* **do**
4:       **case** $1, 2, 5, 6$
5:         Send ID_{leader} on $\mathcal{H}$
6:       **case** $3, 4$
7:         Listen on $\mathcal{H}$
8:         **if** $msg = \emptyset$ **then**
9:           $handshake \leftarrow fail$
10:         **else**
11:           $meet \leftarrow msg.[2]$
12:     **if** $handshake = fail$ **then**
13:       $count \leftarrow 0$, $state \leftarrow \mathbb{A}$
14:     **if** $count = 6$ **then**
15:       $count \leftarrow 0$; $state \leftarrow \mathbb{H}$; $game \leftarrow succ$; $lonely \leftarrow 0$
16:   **case** $\mathbb{L}'$
17:     **switch** *count* **do**
18:       **case** $1, 2, 5, 6$
19:         Listen on $\mathcal{H}$
20:         **if** $msg = \emptyset$ **then**
21:           $handshake \leftarrow fail$
22:       **case** $3, 4$
23:         $meet \leftarrow k$
24:         Send $(ID, meet)$ on $\mathcal{H}$
25:     **if** $handshake = fail$ **then**
26:       $count \leftarrow 0$, $state \leftarrow \mathbb{A}$
27:     **if** $count = 6$ **then**
28:       $count \leftarrow 0$; $state \leftarrow \mathbb{L}$; $game \leftarrow succ$; $lonely \leftarrow 0$

Handshake and Red-Blue Protocol. Pseudo-code for these two protocols can be found in Algorithms 5 and 6. In short, a node $h \in \mathbb{H}'$ that received a message in the herald protocol sends for two rounds on the *handshake channel* $\mathcal{H}$, then listens twice, and sends again for two rounds. A node $l \in \mathbb{L}'$ that was sending before acts reversely, i.e., it listens, sends and listens. Only if a node receives *all* expected messages it moves forward to the red-blue protocol, otherwise it returns to $\mathbb{A}$. The handshake can only possibly be completed if a pair of exactly one broadcaster and one receiver participates.

The *red-blue protocol* is a repetition of $\tau_{\textit{red-blue}}/8 = \Theta(\log n)$ *red-blue games* of 8 rounds each. In odd rounds, both nodes l and h of the leader-herald pair send a blocking signal on $\mathcal{H}$, preventing nearby handshakes to succeed. At the beginning of each game, the leader l picks randomly blue or red. If it picked red, then in round 2 it will send a message on channel $\mathcal{G}$ and listens on $\mathcal{G}$ in round 4, for blue it acts reversely. In round 6, l sends on a previously decided meeting channel $\mathcal{R}_k$ the index k' of the meeting channel for the next red-blue game.[4] In round 8 it listens on $\mathcal{R}_{k'}$. h on the other hand sends a message in both rounds 2 and 4. It listens in round 6 to update the meeting channel and in round 8 it sends a message on $\mathcal{R}_{k'}$.

By design of the handshake and the blocking signals of odd rounds in the red-blue protocol, a leader l can neighbor a leader or herald of a different pair *only* if that other node moved to the red-blue protocol simultaneously or with a 2-round shift. If l does have such a neighbor, at some point it will not hear its herald in round 2/4, when it listens. l then aborts the red-blue protocol, notifies h in round 6 and returns to $\mathbb{A}$. The messages sent by l/h in round 6/8 also have the purpose of letting nearby listening active nodes

[3] We want to note that π_ℓ is a constant parameter that we can choose arbitrarily.

[4] The very first meeting channel is fixed by l during the handshake, confer lines 11 and 24 of Algorithm 5.

Algorithm 6 RedBlueGame

1: **switch** *state* **do**
2: **case** $\mathbb{H}$
3: $\gamma \leftarrow \max\left\{\gamma\sigma_{\oplus}^{-20}, \gamma_{min}\right\}$
4: **switch** *count* **do**
5: **case** $1,3,5,7 \mod 8$ ▷ block $\mathcal{H}$
6: Send $(state, ID_{leader})$ on $\mathcal{H}$
7: **case** $2 \mod 8$ ▷ help leader with game
8: Send $(state, ID_{leader})$ on $\mathcal{G}$
9: **case** $4 \mod 8$ ▷ help leader with game
10: Send $(state, ID_{leader})$ on $\mathcal{G}$
11: **case** $6 \mod 8$
12: Listen on $\mathcal{R}_{meet}$ ▷ from previous game
13: **if** $msg \neq (ID_{leader}, succ, *)$ **then**
14: $count \leftarrow 0$, $state \leftarrow \mathbb{A}$; $lonely \leftarrow 0$
15: **else**
16: $meet \leftarrow msg.[3]$
17: **if** $count > \tau_{red\text{-}blue}$ **then**
18: $state \leftarrow \mathbb{E}$
19: **case** $8 \mod 8$ ▷ notify neighbors
20: Send $(state, ID_{leader})$ on $\mathcal{R}_{meet}$
21: **case** $\mathbb{L}$
22: $\gamma \leftarrow \max\left\{\gamma\sigma_{\oplus}^{-20}, \gamma_{min}\right\}$
23: **switch** *count* **do**
24: **case** $1,3,5,7 \mod 8$ ▷ block $\mathcal{H}$
25: **if** $count \pmod 8 = 1$ **then**
26: pick randomly $color \in \{red, blue\}$
27: Send $(state, ID)$ on $\mathcal{H}$
28: **case** $2 \mod 8$ ▷ red-blue game
29: **if** $color = blue$ **then**
30: Listen on $\mathcal{G}$;
31: **if** $msg = \emptyset$ or $ID \notin msg$ **then**
32: $game \leftarrow fail$
33: **else** Send (ID) on $\mathcal{G}$
34: **case** $4 \mod 8$ ▷ red-blue game
35: **if** $color = red$ **then**
36: Listen on $\mathcal{G}$;
37: **if** $msg = \emptyset$ or $ID \notin msg$ **then**
38: $game \leftarrow fail$
39: **else** Send (ID) on $\mathcal{G}$
40: **case** $6 \mod 8$ ▷ Send *game* & new $\mathcal{R}_{meet}$
41: Send $(ID_{leader}, game, k)$ on $\mathcal{R}_{meet}$
42: $meet \leftarrow k$
43: **if** $game = fail$ **then**
44: $count \leftarrow 0$, $state \leftarrow \mathbb{A}$; $lonely \leftarrow 0$
45: **if** $count > \tau_{red\text{-}blue}$ **then**
46: $state \leftarrow \mathbb{M}$
47: **case** $8 \mod 8$
48: Listen on $\mathcal{R}_{meet}$

reduce their activity values. An *isolated pair* on the other hand cannot be knocked out anymore[5] and after $\tau_{red\text{-}blue} = \Theta(\log n)$ rounds the pair can assume that w.h.p. there is no other conflicting pair nearby and the leader joins the MIS.

The handshake did not change and the red-blue protocol has been extended by 2 rounds—now heralds also can reach their neighbors, confer lines 19, 20, 47, 48 of Algorithm 6. Rounds 1-6 are untouched. Unlike in [9], each round a node spends in the red-blue game, it decreases its activity significantly—after all it is getting messages from a leader or herald all the time. This is accounted for in lines 3 and 22.

Summary of Changes. Compared to the algorithm in [9], the following three things have changed. The *loneliness support block* is not executed anymore, except for maintaining the counter *lonely*. Also, the threshold τ_{lonely} has been lowered to $\Theta\left(\log^2 n/\mathcal{F} + \log n\right)$ to reflect the new runtime of the algorithm. The main change is that nodes reduce their activity γ if they hear from a nearby leader or herald. The change in the red-blue game is an addition of 2 rounds: the seventh round is just a copy of rounds 1, 3 and 5; the eighth round gives the herald of the pair a possibility to notify nearby active nodes in order to reduce their activity values—so far only leaders and MIS nodes were able to reach out to their neighbors.

Note that while the algorithm itself has barely changed, the analysis needed to be extended vastly to reduce the runtime of the algorithm to optimal values.

[5] except by an MIS node, but that already implies progress

4 Analysis

4.1 Approach

To prove that Algorithm 1 indeed solves MIS in the given time bounds, we take the following approach. In [9] it was proven that the graph, induced by nodes that passed through the decay filter, has maximum degree $\Delta_{max} = O(\log^3 n)$. A node u in the herald filter ($u \in V_{\text{hf}}$) enters the MIS either if it assumes to be alone, or if it manages to create and maintain a leader-herald bond with a neighboring node for $\tau_{\textit{red-blue}} = \Theta(\log n)$ rounds. Once $u \in V_{\text{hf}}$, it either enters the MIS due to assumed loneliness; or if u has a neighbor in V_{hf}, then within radius $\delta := \delta_\alpha = \Theta(\log \log n)$ soon a leader-herald pair is created that maintains its bond for $\tau_{\textit{red-blue}}$ rounds.[6] So far this is the same as in [9]. There, however, a stagnation of up to $\tau_{\textit{red-blue}}$ rounds might follow before the next isolated leader-herald pair or MIS node gets created in $N^\delta(u)$. Considering that up to $\alpha(\delta)$ nodes in $N^\delta(u)$ can enter the MIS before u or one of its neighbors enters itself, the runtime of the herald filter is $O(\tau_{\textit{red-blue}}\alpha(\delta))$, or $O(\log n \operatorname{polyloglog} n)$ if α is polynomial.

In the present paper, by decreasing activity levels of nodes neighboring leader-herald pairs, the stagnation that can be caused by leaders on their way to join the MIS does not last for longer than $O(\log \log n)$ rounds in expectation. This allows the creation of isolated leader-herald pairs in $N^\delta(u)$ in a pipelined manner, reducing the expected runtime of the herald filter to $O(\alpha(\delta) \log \log n) = O(\alpha^\delta \log \log n)$. Unlike in [9], here we also can choose δ as an arbitrarily small value in $\Theta(\log \log n)$ without increasing the runtime by more than constant factors. Choosing $\delta < \log \log n / \log \alpha$ and a Chernoff argument bounds the runtime of the herald filter by $O(\log n)$ with high probability.

In more detail, let u be a node that enters the herald filter in round t_u. For the sake of contradiction assume that u is not decided by time $t_u + \tau_{\textit{runtime}}$. If u stays lonely, it enters the MIS eventually in $\tau_{\textit{lonely}} \ll \tau_{\textit{runtime}}$ rounds. Note that for u to move from being non-lonely to lonely, some node in $N^2(u)$ must have entered the MIS shortly before that and eliminated all neighbors that u had in V_{hf}. This can happen at most α^2 times and thus the time u spends lonely is at most $\alpha^2 \tau_{\textit{lonely}} \ll \tau_{\textit{runtime}}$. Hence, assume that u is not lonely, i.e., has a neighbor u', and that no node in $N^2(u)$ joins the MIS. We show that then most of the time both u and u' have a high activity value γ.

The following argumentation motivates this. For a node u to decrease $\gamma(u)$, it must neighbor a pair. Let us call isolated pairs (in which the leader does not neighbor another leader or herald) *good pairs* and the others *bad pairs*. Conditioning on the event of a pair being created, there is a constant probability that it is a good pair. This can be considered progress, as it guarantees one of two things: Within $O(\log n)$ rounds either the leader of the good pair itself enters the MIS or a neighbor of this pair does. In the opposite case of bad pairs being created, in expectation these remain bad pairs only for a constant number of rounds. Moreover, w.h.p., there are no more than $O(\log n)$ rounds *in total* in which bad pairs exist in $N^3(u)$ after t_u, also causing at most $O(\log n)$ rounds of u and u' having an activity value below $1/2$. Adjusting parameters we get that for some $\tau_{\textit{progress}} = O(\tau_{\textit{lonely}})$ and an arbitrarily small constant ε, for $(1-\varepsilon)\tau_{\textit{progress}}$ rounds in $[t_u, t_u + \tau_{\textit{progress}}]$ the activity values of both u and u' are $1/2$.

Furthermore, all pairs, good and bad, inform their neighbors. By the definition of good pairs, the leaders of these form an independent set. With our choice of δ thus at most $O(\sqrt{\log n})$ good pairs exist in $N^\delta(u)$. We argue that the activity values of nodes neighboring a pair that participated in the red-blue protocol for $\Omega(\log \log n)$ rounds (which almost surely holds for good pairs), are below $\gamma_{\textit{low}} := \Theta(1/\operatorname{polylog} n)$ with some decent probability (i.e., $1 - \log^{\Omega(1)} n$). The total number of nodes in $N^\delta(u)$ becoming part of a good pair in $[t_u, t_u + \tau_{\textit{progress}}]$ is $O(\sqrt{\log n})$ and hence the total amount of nodes neighboring good pairs in that time is $O(\sqrt{\log n}\Delta_{max})$. A union bound and a Chernoff bound provide that the total amount of rounds in which any node v in $N^\delta(u)$ neighboring a good pair has an $\gamma(v) > \gamma_{\textit{low}}$ is less than $\varepsilon\tau_{\textit{progress}}$.

Together with the previous claim we get that in $(1-2\varepsilon)\tau_{\textit{progress}}$ rounds in $[t_u, t_u + \tau_{\textit{progress}}]$ both con-

[6]"Soon" indeed means in $O(1)$ rounds in expectation, as long as $\mathcal{F} = \Omega(\log \log n)$.

ditions are true: $\gamma(u) = \gamma(u') = 1/2$ and *all* good pairs in $N^\delta(u)$ "silenced" their neighbors—i.e., all their neighbors have activity below γ_{low}. Let us call a round with this property *promising for* u. Without going into detail, we can show that now within distance δ there exists a node w with the property of being so-called η-fat, i.e., w's neighborhood is at least roughly as active as that of any of its neighbors'. Fatness implies that w.c.p. two nodes l and h in $N^1(w)$ become a *good* leader-herald pair. As said before, such a pair reduces the activity values of its neighbors rather quickly, which causes the property of η-fatness to move away from w to another node in $N^\delta(u)$ and we can repeat the argument. If a bad pair is created, then η-fatness might shortly fade, but is restored quickly, so we can almost omit this case. Again using Chernoff tail bounds, we show that at some point u itself becomes η-fat and now the creation of an MIS node in $N^2(u)$ is inevitably.

We summarize again. Once an MIS node or good pair arises in constant distance from u, we are done. In an $\Omega(\log^2 n/\mathcal{F} + \log n)$ interval, u is mostly in a promising state. W.c.p. every $O(1)$ rounds a node in $N^\delta(u)$ becomes part of a good pair or joins the MIS. In expectation, within $\Theta(\log\log n)$ rounds MIS nodes eliminate their neighbors and good pairs *silence* theirs. After any of those events happen, we measure the time until u is in a promising state again. Using Chernoff over $O(\sqrt{\log n})$ such random variables results in needing at most $O(\log n)$ time, thus, by then u must be covered.

4.2 Guarantees from the Decay Filter

We informally state the two main accomplishments of the decay filter, proper lemma statements are below; for proofs we refer to [9]. For each node v the decay filter guarantees that within the runtime of $\Theta(\log^2 n/\mathcal{F} + \log n)$ rounds,

(1) v or one of its neighbors enters the herald filter, but

(2) no more than $\Delta_{max} = O(\log^3 n)$ nodes in $N^1(v)$ do.[7]

From now on we only look at the graph G' induced by $V' := V_{\text{hf}} \cup \mathbb{M}$, induced by non-eliminated nodes that made it past the decay filter. *All* notations are tied to this subgraph, though we omit this in our notations, i.e., $N(u)$ means the neighborhood of u in G'. Instead, if we need to consider nodes from the states $\mathbb{W}$, $\mathbb{D}$, then we explicitly say so and show this e.g. by writing $N_G(u)$.

Lemma 4.1. *W.h.p., for each node v and each round r, at most $O(\log n)$ nodes in $N^1_G(v)$ come out of the decay filter in round r to enter the herald filter. Each node that enters the herald filter has spent $\Omega(\log n)$ rounds in the decay filter.*

Lemma 4.2. *W.h.p., for each node u that is in the decay filter in round r, by round $r' = r + O(\log^2 n/\mathcal{F} + \log n)$, either u is* dominated, *in which case it has a dominating neighbor, or at least one node in $N^1(u)$ gets out of the decay filter and enters the herald filter.*

This statement is the same as Lemma 8.4 in [9], except that there the bound was listed as $r' = r + O(\log^2 n/\mathcal{F}) + \tilde{O}(\log n)$. Yet the proof in [9] does actually already support the bound $r' = r + O(\log^2 n/\mathcal{F} + \log n)$. In Algorithm 2, DFILTER, we changed the style of the algorithm compared to the one in [9], but not the way the algorithm works, hence we omit the proof for Lemma 4.2 and refer to [9].[8]

4.3 Definitions for the Herald Filter

Practically all parameters (including the above mentioned Δ_{max}) depend in one way or another on the bound on independence, i.e., on α, but in most cases those dependencies are captured in the hidden constants of those asymptotic bounds.

For our analysis of a node u that enters the herald filter, we observe a specific $\delta = \Theta(\log\log n)$ neighborhood $N^\delta(u)$ of u. We set

[7] For large enough n it holds that $\Delta_{max} \le \log^4 n$ and we assume this in our analysis.

[8] The underlying algorithm has been first used and analyzed in [11], in a slightly more restrictive graph model and in [8] it was shown that it also works in BIGs.

$$\delta := \delta_\alpha := \frac{\log\log n}{2\log\alpha} = \Theta(\log\log n), \tag{1}$$

i.e., $\alpha^\delta = (2^{\log\alpha})^{\frac{\log\log n}{2\log\alpha}} = \sqrt{\log n}$. The choice of δ guarantees that any independent set in a δ-neighborhood is of size at most $\sqrt{\log n}$.

Our main goal is to show quick progress in $N^\delta(u)$. Progress is clearly achieved if an MIS node arises, but due to the way a node can become an MIS node, we also consider the creation of an isolated leader-herald pair progress (more precisely, the *leader* of the pair needs to be isolated from other nodes in $\mathbb{L}$ or $\mathbb{H}$), as the leader will eventually join the MIS (or be knocked out permanently by a newly created MIS node).

Definition 4.3. ***(Good Pair, Bad Pair)*** *Consider a leader-herald pair (l, h) in round r. We say (l, h) is a* good pair *in round r if none of the neighbors of l (other than h) is (1) in state $\mathbb{L}$ or (2) in state $\mathbb{H}$ or (3) is a herald candidate in round 5 or 6 of its respective handshake. Otherwise we say that (l, h) is a* bad pair.

Note that the definition of a good/bad pair is independent of possibly neighboring MIS nodes. MIS nodes existing already for 4 rounds prevent the creation of leader-herald pairs in their neighborhood completely. If on the other hand a new MIS node appears next to a leader-herald pair (which is w.h.p. only possible through the loneliness route), then we have progress in a close neighborhood. Also, note that only the leader of the pair must be 'isolated'. There are two reasons for this: (1) only leaders join the MIS (2) by protocol design the *herald* of a pair can only receive messages from MIS nodes or its own leader—not by other leaders (not even in round 6) nor other heralds. This is due to the fact that any neighboring heralds act completely synchronously and a leader neighboring a non-paired herald is ahead by precisely 2 rounds.[9] Note also that bad pairs can become good, but not vice versa. This is because all leaders and heralds prevent the creation of further leaders/heralds in their neighborhood.

Definition 4.4. ***(Activity Mass)*** *For a node u we define $\Gamma(u) := \sum_{v\in N^1(u)} \gamma(v)$. We call this the* activity sum *or* activity mass *of node u. Furthermore we let $\Gamma^\circ(u) := \Gamma(u) - \gamma(u) = \sum_{v\in N(u)} \gamma(v)$. In some cases we are only interested in the activity mass of active nodes and then we have $\Gamma_{\mathbb{A}}(u) := \sum_{v\in N^1_{\mathbb{A}}(u)} \gamma(v)$ and $\Gamma^\circ_{\mathbb{A}}(u)$ is defined analogously. Also*

$$\gamma_{min} := \log^{-24} n = \Theta(1/\operatorname{polylog} n), \tag{2}$$

$$\gamma_{low} := \sqrt{\gamma_{min}} = \log^{-12} n. \tag{3}$$

Definition 4.5. ***(Fatness)*** *We call a node u $\hat\eta$-*fat *for some $\hat\eta \in (0,1)$, if $\Gamma(u) \geq \hat\eta \cdot \max_{v\in N(u)}\{\Gamma(v)\}$.*

In simple words, in terms of activity mass, u is (at least) in the same 'league' as its neighboring nodes. Using this we choose a specific fatness parameter $\eta < 1$:

$$\eta = \eta_\alpha := \alpha^{-8} \leq \alpha^{-2\frac{\log\Delta_{max}}{\log\log n}} \tag{4}$$

The choice of η assures that a chain of activity sums $(\Gamma(v_i))_{i\geq 1}$ of nodes v_i on a path $v_1, v_2, v_3, \ldots$ with $\Gamma(v_i) \geq \eta^{-1}\Gamma(v_{i-1})$ and $\Gamma(v_1) \geq 2$ has length at most δ, because

$$(\eta^{-1})^\delta = (\alpha^{-8})^{-\frac{\log\log n}{2\log\alpha}} \geq (2^{\log\alpha})^{2\frac{\log\Delta_{max}}{\log\log n}\frac{\log\log n}{2\log\alpha}} = \Delta_{max} \overset{\gamma(u)\leq 1/2}{>} \max_{u\in V_{\text{hf}}} \Gamma(u). \tag{5}$$

The algorithm needs to know a few more parameters. $\sigma_\oplus$ and $\sigma_\ominus$ govern the changes in a node's activity level. The former is a small constant, greater than, but close to 1. In most rounds a node u increases $\gamma(u)$ by

[9] Cf. Lemma 4.10 and Algorithm 6.

$\sigma_\oplus$. $\sigma_\ominus$ is a much larger factor used for decreasing activity, large enough to undo many previous increments, but still in $O(1)$.

$$\sigma_\oplus \quad := \quad 2^{6/(1000\bar{m})} > 1 \tag{6}$$
$$\sigma_\ominus \quad := \quad \sigma_\oplus^{20\bar{m}} = 2^{12/100} > 1 \tag{7}$$

$\bar{m}$ is a large enough constant that depends on $n_\mathcal{R}$, but assuming that $n_\mathcal{R} \geq 3\alpha^2$, $\bar{m} \geq 2^{16} n_\mathcal{R}$ suffices. Since $\gamma_{\mathit{min}} = \log^{-24} n$, $167\bar{m}\log\log n = \Theta(\log\log n)$ consecutive increments raise a node's activity value to $1/2$. Analogously, $\Theta(\log\log n)$ decrements decrease it to its minimal value γ_{min}.

Also two time thresholds $\tau_{\mathit{red\text{-}blue}} = \Theta(\log n)$ and $\tau_{\mathit{lonely}} = \Theta(\log^2 n/\mathcal{F} + \log n)$ are needed by the algorithm. $\tau_{\mathit{red\text{-}blue}}$ is the number of rounds a node spends in the red-blue protocol, and it is a multiple of 8. If a node $u \in V_{\text{hf}}$ does not receive a single message for τ_{lonely} consecutive rounds, while being in the herald filter, a u deduces that it is alone or all its neighbors got eliminated, and joins the MIS. In our analysis we use further time thresholds $\tau_{\mathit{notification}} = \Theta(\log n)$, $\tau_{\mathit{progress}} = \Theta(\log^2 n/\mathcal{F} + \log n)$ and $\tau_{\mathit{runtime}} = \Theta(\log^2 n/\mathcal{F} + \log n)$, for which the following inequality chain holds:

$$\tau_{\mathit{runtime}} \gg \tau_{\mathit{lonely}} \gg \tau_{\mathit{progress}} \gg \tau_{\mathit{red\text{-}blue}} \gg \tau_{\mathit{notification}}$$

$\tau_{\mathit{notification}}$ is the maximum time needed for an MIS node to notify, w.h.p., all its neighbors. If a node u is not lonely, then, w.h.p., significant progress is achieved in less than τ_{progress} rounds; more precisely, an MIS node is created in $N^{O(1)}(u)$. W.h.p., τ_{runtime} is the maximum time a node spends in the herald filter before it gets decided.

4.4 Candidate Election—Nodes in States $\mathbb{A}$ (and $\mathbb{L}'$)

At first we establish a few facts about how nodes can transit from state $\mathbb{A}$ to state $\mathbb{L}'$ or $\mathbb{H}'$, respectively. Note that nodes can switch between states $\mathbb{A}$ and $\mathbb{L}'$ without communication, but to get towards any of the three states $\mathbb{H}'$, $\mathbb{L}$ and $\mathbb{H}$, communication is mandatory.

The next lemma contains a variety of events. To not clutter the lemma statement, we define them here. Let k be a positive integer *constant*, r some round, u some node in the herald filter in round r, i an index from $1, \ldots, n_\mathcal{A}$, S be a (possibly empty) subset of $N^k(u) \cap \mathbb{A}$. Furthermore let $\partial S \subset S$ be *the* subset of S that has connections outside of S, but in $N^k(u)$, i.e., $\partial S := S \cap N\big(N^k(u) \setminus S\big)$. At last, let $S = S^n \dot\cup S^b \dot\cup S^l$ be a partition of S. We call the tuple (k, u, r, i, S)) a *constellation*. For a constellation the following events are defined:

- $\mathcal{S}_{\neg i}/\mathcal{S}^n_{\neg i}$: *no* node in S/S^n operates on $\mathcal{A}_i$ in round r,
- $\mathcal{S}_i/\mathcal{S}^b_i/\mathcal{S}^l_i$: *all* nodes in $S/S^b/S^l$ operate on $\mathcal{A}_i$ in round r,
- $\partial\mathcal{S}_{\neg i}$: *no* node in ∂S operates on $\mathcal{A}_i$ in round r,
- H_i: *no* node in $N^k(u) \setminus S$ receives a message on channel $\mathcal{A}_i$ in round r,
- $H_{\neg i}$: *no* node in $N^k(u)$ receives a message on some channel $\mathcal{A}_j \neq \mathcal{A}_i$ in round r and
- H: *no* node in $N^k(u)$ receives a message on any channel in $\{\mathcal{A}_1, \ldots, \mathcal{A}_{n_\mathcal{A}}\}$ in round r.

Lemma 4.6. *Let (k, u, r, i, S) be a constellation. Then,*

(1) $\mathbf{P}(H) = 1 - O(\pi_\ell \alpha^k)$

(2) $\mathbf{P}(H_{\neg i} | \mathcal{S}_{\neg i}) = 1 - O(\pi_\ell \alpha^k)$

(3) $\mathbf{P}(H_{\neg i} | \mathcal{S}_i) = 1 - O(\pi_\ell \alpha^k)$

(4) $\mathbf{P}(H_{\neg i} | \mathcal{S}^n_{\neg i} \wedge \mathcal{S}^b_i \wedge \mathcal{S}^l_i) = 1 - O(\pi_\ell \alpha^k)$

(5) $\mathbf{P}(H_i | \partial\mathcal{S}^n_{\neg i}) = 1 - O(\pi_\ell \alpha^k)$

The proofs for these statements are provided in Appendix B.1.

Look at (1). The lemma says that the probability for a herald candidate to be created in any single round for any neighborhood of constant radius is at most linear in π_ℓ. Since π_ℓ is an arbitrarily small constant parameter chosen by us, we can make the probability for this event arbitrarily small. The proof for (1) is exactly the same as the proof for Lemma 8.6 in [8], with $\alpha(k)$ replaced by α^k. We provide its proof in Appendix B.1 nevertheless, as (2) and (3) are new results that do directly depend on (1). What (2) says, is, that even if we condition on some nodes $S \subset N^k(u)$ *not* to operate on channel $\mathcal{A}_i$, this does not increase the chance (significantly—i.e., by more than a constant factor) for these or other nodes to receive anything on some channel $\mathcal{A}_j \neq \mathcal{A}_i$. Analogously, (3) claims that if we condition on some nodes $S \subset N^k(u)$ to *definitely* operate on channel $\mathcal{A}_i$, this does still not affect the probability for any other node to receive any message on some other channel $\mathcal{A}_j \neq \mathcal{A}_i$. (4) is a combination of (2) and (3), plus we even fix the knowledge of which nodes, that operate on $\mathcal{A}_i$, do broadcast (S^b) or listen (S^l). (5) has already been proven in [8] as Claim 8.9, except that there k was fixed to 2. The analogous proof is provided in the Appendix as well.

Under certain conditions the creation of herald candidates can be lower bounded. However, for our algorithm to work, we need not only to prove that they are created, but that this creation happens in *solitude*, i.e., in a close neighborhood no other herald candidates are created. Hence the next lemma is a key result in our whole proof. It is almost Lemma 8.8 from [8], however with some adaptions. To state the lemma, we need to introduce the parameters

$$\gamma_{\mathit{min}} := \log^{-6\kappa_\Delta} n = \frac{1}{\Omega(\operatorname{polylog} n)} \quad (8)$$

$$\gamma_{\mathit{low}} := \sqrt{\gamma_{\mathit{min}}} = \log^{-3\kappa_\Delta} n, \quad (9)$$

depending on the constant $\kappa_\Delta \geq \log_{\log n} \Delta_{\mathit{max}}$.

Lemma 4.7. *Let t be a round in which for a node u in state $\mathbb{A}$ in the herald filter the following holds:*

- *there is no herald candidate in $N^2(u)$,*
- *all nodes $v \in N^2(u)$ that neighbor a herald or leader, have $\gamma(v) \leq \gamma_{\mathit{low}}$,*
- *all nodes in $N^2(u)$ neighboring MIS nodes are eliminated,*
- $\Gamma(u) \geq 1$,

If in addition it holds that either
(a) $\Gamma(u) < 5\alpha$, *u is $\frac{1}{5\alpha}$-fat and $\gamma(u) = \frac{1}{2}$, or*
(b) $\Gamma(u) \geq 5\alpha$ *and u is η-fat.*
Then by the end of round $t' \in [t, t+7]$, with probability $\Omega(\pi_\ell)$ either a node in $N^2(u)$ joins the MIS or a good pair $(l, h) \in (\mathbb{L} \cap N^1(u)) \times (\mathbb{H} \cap N^1(u))$ is created.

Let us start with an intuition of this Lemma. The basic intention is to show that if u is η-fat, then w.c.p. in constant many rounds a *good* leader-herald pair with both endpoints in $N^1(u)$ arises—for this u itself does not have to have a high activity value, i.e., u does not need to be a likely part of the leader-herald pair. The lemma lists many requirements. We show later that shortly after a node v moves to the herald filter, within distance $\delta = O(\log \log n)$ most of the time there exists a node u that satisfies these conditions. We also show that if an isolated pair is created in $N^1(u)$, those requirements are again satisfied $O(\operatorname{polyloglog} n)$ rounds later (in expectation) by another node u' in this δ-neighborhood of v.

The proof follows mostly the lines of the one in [9], but there a constant number of factors of $1/2$ accumulate, while here they have to be exchanged by factors of η—but this only causes changes in asymptotically irrelevant constants. In addition, nodes neighboring leaders or heralds need special attention. A full detailed proof of this new version is listed in Appendix B.2. There are few things we want to elaborate.

Within the neighborhood of a fat node u with activity mass at least one, w.c.p. "good things" happen (i.e., the creation of MIS nodes or good leader-herald pairs) within constant many rounds, *even if* there are herald candidates nearby or if some nodes neighboring *bad* leaders or heralds have high activity values. In other words, the first two statements could be relaxed a bit. Instead we use other results to show that from those relaxed conditions one can get to the tighter conditions listed here w.c.p. in constant many rounds. Also note that we allow 7 rounds to pass, even though the handshake only needs 6 rounds. This is due to the fact that we require a certain property that might not be given in round t', but, w.c.p., is given one round later. As a last remark we want to say that 5α in this lemma could be replaced by any constant greater than $\frac{20}{7}\alpha$ and η could be replaced by any fatness constant smaller than 1.

We argue in the subsections about the handshake protocol and the red-blue protocol, Subsections 4.5 and 4.6, that every time an isolated pair is created, the algorithm achieves progress, as it guarantees the creation of an MIS node nearby—even if this event is delayed by $O(\log n)$ rounds.

Therefore, Lemma 4.7 "promises" progress in the proximity of a fat node. However, we do not have such a statement for areas without fat nodes. As we describe a bit more in detail later, an excessive creation of bad pairs in such areas can even cause problems for our argumentation. However, the next result shows that *if* a pair is created at all, then w.c.p. this pair is good. This allows us to proof later in Lemma 4.14 that nodes in the herald filter are practically always very active in the candidate election process—unless they already neighbor an MIS node or a good pair.

Lemma 4.8. *Let r be a round in which node u is in state $\mathbb{A}$ and $N_{\mathbb{A}}(u) \neq \emptyset$. Let B^u be the event that at the end of round r, u moves to state $\mathbb{H}'$ due to receiving a message from some node $v \in N_{\mathbb{A}}(u)$ on some channel $\mathcal{A}_{\bar{\lambda}}$. Further, let $D^u \subseteq B^u$ be the event that B^u holds and in addition no other node $v' \in N^3(v) \setminus \{u\}$ receives any message on channel $\mathcal{A}_{\bar{\lambda}}$ in round r. It holds that*

$$\mathbf{P}(B^u) \quad = \quad \begin{cases} O\left(\pi_\ell \frac{\gamma(u)}{\Gamma^\circ_{\mathbb{A}}(u)}\right) & \Gamma^\circ_{\mathbb{A}}(u) > 2 \\ O\left(\pi_\ell \gamma(u) \Gamma^\circ_{\mathbb{A}}(u)\right) & \Gamma^\circ_{\mathbb{A}}(u) \leq 2 \end{cases}, \tag{10}$$

$$\mathbf{P}(D^u) \quad = \quad \begin{cases} \Omega\left(\pi_\ell \frac{\gamma(u)}{\Gamma^\circ_{\mathbb{A}}(u)}\right) & \Gamma^\circ_{\mathbb{A}}(u) > 2 \\ \Omega\left(\pi_\ell \gamma(u) \Gamma^\circ_{\mathbb{A}}(u)\right) & \Gamma^\circ_{\mathbb{A}}(u) \leq 2 \end{cases}. \tag{11}$$

Proof. In the calculations below we make use of the following inequalities.

$$1 - \pi_\ell \quad \geq \quad 0.9 \tag{12}$$

$$1 - x \quad \leq \quad e^{-x}, \quad \forall x \tag{13}$$

$$1 - x \quad \geq \quad e^{-2x}, \quad \forall x \in [0, 1/2] \tag{14}$$

$$\prod_{w \in A \setminus B} f(w) \quad \leq \quad F^{|B|} \prod_{w \in A} f(w), \quad \text{if } f(w) \geq F^{-1} \tag{15}$$

If $\Gamma^\circ_{\mathbb{A}}(u) > 2$, for simplicity, we assume that $\log \Gamma^\circ_{\mathbb{A}}(u)$ is a positive integer. It becomes clear from the proof that for non-integer values an adaption is straightforward, but hard to read.

We let $B_i^{u,v}$ be the event that an *active* node u receives a message from an *active* node v on channel $\mathcal{A}_i$, i.e., u listens and v broadcasts on $\mathcal{A}_i$, while no other node $w \in (N_{\mathbb{A}}(u)) \setminus \{v\}$ broadcasts on $\mathcal{A}_i$. For different v these events are disjoint, and we can define $B_i^u := \bigcup_{v \in N_{\mathbb{A}}(u)} B_i^{u,v}$ and we see that $B^u = \bigcup_{i=1}^{n_{\mathcal{A}}} B_i^u$.

More restrictive are the analogously defined events $D_i^{u,v}$, D_i^u, in which we also require that no node $x \in N^3(v) \setminus \{u\}$ receives a message on channel i. In that case, $D^u = \bigcup_{i=1}^{n_{\mathcal{A}}} D_i^u$ is the event that nodes u and v engage in the handshake protocol in the round after they met on some channel $\mathcal{A}_{\bar{\lambda}}$, without other nodes nearby having received something on that same channel. In particular, no other herald candidates try to engage with v in the handshake protocol.

Upper Bounds

Let q_i^w be the probability that node w does broadcast on channel $\mathcal{A}_i$, i.e., $q_i^w = (1-\pi_\ell)\gamma(w)2^{-i} \le 1/4$, and accordingly $\bar{q}_i^w := 1 - q_\lambda^w \ge 3/4$ is the probability that w does not broadcast on channel $\mathcal{A}_i$. We denote with p_i^w the probability that w listens on channel $\mathcal{A}_i$.

$$\begin{aligned}\mathbf{P}(B_i^{u,v}) &= p_i^u q_i^v \prod_{w \in N_{\mathbb{A}}(u)\setminus\{v\}} \bar{q}_i^w = \pi_\ell \gamma(u) 2^{-i} (1-\pi_\ell)\gamma(v)2^{-i} \prod_{w \in N_{\mathbb{A}}(u)\setminus\{v\}} \left(1-(1-\pi_\ell)\gamma(w)2^{-i}\right) \\ &\overset{(13),(15)}{\le} \pi_\ell \gamma(u)\gamma(v) 2^{-2i} \frac{4}{3} e^{-\frac{1}{2}\Gamma_{\mathbb{A}}^{\circ}(u)2^{-i}}\end{aligned}$$

$$\mathbf{P}(B_i^u) \le \pi_\ell \gamma(u) \Gamma_{\mathbb{A}}^{\circ}(u) 2^{-2i+1} e^{-\Gamma_{\mathbb{A}}^{\circ}(u)2^{-i-1}} = 8\pi_\ell \frac{\gamma(u)}{\Gamma_{\mathbb{A}}^{\circ}(u)} (\Gamma_{\mathbb{A}}^{\circ}(u)2^{-i-1})^2 e^{-\Gamma_{\mathbb{A}}^{\circ}(u)2^{-i-1}} =: C_i^u \tag{16}$$

Consider the case $\Gamma_{\mathbb{A}}^{\circ}(u) > 2$. For any $\zeta > 0$ it holds that $\zeta^2 e^{-\zeta} = O(1)$, and by using $\zeta = \Gamma_{\mathbb{A}}^{\circ}(u)2^{-i}$ and we get that for any fixed i

$$C_i^u = O\Big(\pi_\ell \frac{\gamma(u)}{\Gamma_{\mathbb{A}}^{\circ}(u)}\Big) \quad \text{and} \quad C_i^u = \Theta(C_{i+1}^u) \tag{17}$$

Furthermore, with $\lambda := \log \Gamma_{\mathbb{A}}^{\circ}(u)$:

$$\begin{aligned} \frac{C_{i+1}^u}{C_i^u} &= \frac{1}{4} e^{\frac{1}{4}\Gamma_{\mathbb{A}}^{\circ}(u)2^{-i}} < \frac{1}{2} & \forall i \ge \lambda \\ \frac{C_i^u}{C_{i+1}^u} &= 4 e^{-\frac{1}{4}\Gamma_{\mathbb{A}}^{\circ}(u)2^{-i}} < \frac{1}{2} & \forall i \le \lambda - 4 \end{aligned}$$

Thus, the sum of all C_i^u with i larger than λ can be upper bounded by C_λ^u using the geometric series, and, if $\lambda > 4$, all $i \le \lambda - 4$ can be upper bounded by $C_{\lambda-4}^u$. However, due to (17), $C_{\lambda-4}^u + \cdots + C_\lambda^u$ is in $O(C_\lambda^u) = O\big(\pi_\ell \frac{\gamma(u)}{\Gamma_{\mathbb{A}}^{\circ}(u)}\big)$. In total we get that $\mathbf{P}(B^u) = O\big(\pi_\ell \frac{\gamma(u)}{\Gamma_{\mathbb{A}}^{\circ}(u)}\big)$.

Now consider the case $\Gamma_{\mathbb{A}}^{\circ}(u) \le 2$. In (16) we upper bound $e^{-\Gamma_{\mathbb{A}}^{\circ}(u)2^{-i-1}}$ by 1 and get

$$\mathbf{P}(B_i^u) \le \pi_\ell \gamma(u) \Gamma_{\mathbb{A}}^{\circ}(u) 2^{-2i+1} =: \hat{C}_i^u \quad \text{and} \quad \hat{C}_{i+1}^u \le \frac{1}{4}\hat{C}_i^u.$$

Clearly, $\hat{C}_1^u \le \sum_{i=1}^{n_{\mathcal{A}}} \hat{C}_i^u \le 2\hat{C}_1^u$ due to the convergence of the geometric series. In other words:

$$\mathbf{P}(B^u) \le 2C_1^u = O(\pi_\ell \gamma(u) \Gamma_{\mathbb{A}}^{\circ}(u)).$$

This concludes the proof of the first part of the claim, equations (10). Note that in both cases $\Gamma_{\mathbb{A}}^{\circ}(u) \le 2$ and $\Gamma_{\mathbb{A}}^{\circ}(u) > 2$ we had that

$$\mathbf{P}(B_\lambda^u) = \Omega(B^u). \tag{18}$$

Lower Bounds

For lower bounds we study 2 specific channels, depending whether $\Gamma_{\mathbb{A}}^{\circ}(u)$ is greater than 2 or not. More precisely, we study $\mathcal{A}_\lambda$, where $\lambda := \max\{1, \log \Gamma_{\mathbb{A}}^{\circ}(u)\}$ and show that event D^u happens on this channel with the desired probability. The base idea is to show that w.c.p. no node $x \in N_{\mathbb{A}}^3(u)$ with $\Gamma_{\mathbb{A}}^{\circ}(x) \lesssim \Gamma_{\mathbb{A}}^{\circ}(u)$ is listening on channel $\mathcal{A}_\lambda$ at all, while all nodes x with $\Gamma_{\mathbb{A}}^{\circ}(x) \gtrsim \Gamma_{\mathbb{A}}^{\circ}(u)$ experience a collision on $\mathcal{A}_\lambda$.

Case 1: $\boldsymbol{\Gamma_{\mathbb{A}}^{\circ}(u) > 2, \lambda = \log \Gamma_{\mathbb{A}}^{\circ}(u)}$ and $\boldsymbol{2^{-\lambda} = 1/\Gamma_{\mathbb{A}}^{\circ}(u)}$.

We start with some definitions. We define

$$X_m^{u,v} := \left\{x \in N_{\mathbb{A}}^3(v) \setminus \{u,v\} : 10m\Gamma_{\mathbb{A}}^{\circ}(u) \le \Gamma_{\mathbb{A}}^{\circ}(x) < 10(m+1)\Gamma_{\mathbb{A}}^{\circ}(u)\right\}.$$

Let $F_{\lambda,m}^{u,v}$ be the event that *no* $x \in X_m^{u,v}$ receives *any* message broadcasted on $\mathcal{A}_\lambda$ and let E_λ^x be the event that $x \in \mathbb{A}$ *successfully* receives a message on channel $\mathcal{A}_\lambda$. If $x \in N_{\mathbb{A}}(v) \setminus \{u\}$, then $E_\lambda^x | B_\lambda^{u,v}$ implies that x receives v's message. Further, $p_\lambda^{u,v} = \pi_\ell(1-\pi_\ell)\gamma(u)\gamma(v)2^{-2\lambda}$ is the probability that nodes u and v meet on channel $\mathcal{A}_\lambda$ with u listening and v broadcasting.

Apparently,

$$\mathbf{P}(B_\lambda^{u,v}) = p_\lambda^{u,v} \prod_{w \in N_{\mathbb{A}}(u)\setminus\{v\}} \bar{q}_\lambda^w. \tag{19}$$

We want to calculate $\mathbf{P}(E_\lambda^x | B_\lambda^{u,v})$. For that we need to distinguish between the two cases $x \in N_{\mathbb{A}}(v) \setminus \{u\}$ and $x \notin N_{\mathbb{A}}^1(v)$. For the first case, no other neighbor of x than v is allowed to broadcast, while in the second case exactly one neighbor of x needs to broadcast on $\mathcal{A}_\lambda$.

$$x \in N_{\mathbb{A}}(v)\setminus\{u\}: \quad \mathbf{P}(E_\lambda^x|B_\lambda^{u,v}) = \frac{\mathbf{P}(B_\lambda^{u,v} \cap E_\lambda^x)}{\mathbf{P}(B_\lambda^{u,v})}$$
$$= \frac{1}{\mathbf{P}(B_\lambda^{u,v})} \frac{\mathbf{P}(B_\lambda^{u,v})}{\bar{q}_\lambda^x} \pi_\ell \gamma(x) 2^{-\lambda} \prod_{y \in N_{\mathbb{A}}(x)\setminus\{u,v\}} \bar{q}_\lambda^y \le \frac{4}{3}\pi_\ell \frac{\gamma(x)}{\Gamma_{\mathbb{A}}^{\circ}(u)} \prod_{y \in N_{\mathbb{A}}(x)\setminus\{u,v\}} \bar{q}_\lambda^y \tag{20}$$

$$x \notin N_{\mathbb{A}}^1(v): \quad \mathbf{P}(E_\lambda^x|B_\lambda^{u,v}) = \mathbf{P}(E_\lambda^x|u \text{ listens}) = \pi_\ell\gamma(x)2^{-\lambda} \cdot \sum_{y \in N_{\mathbb{A}}(x)\setminus\{u\}} q_\lambda^y \cdot \prod_{z \in N_{\mathbb{A}}(x)\setminus\{u,y\}} \bar{q}_\lambda^z$$
$$= \pi_\ell\gamma(x)2^{-\lambda} \cdot \sum_{y \in N_{\mathbb{A}}(x)\setminus\{u\}} (1-\pi_\ell)\gamma(y)2^{-\lambda} \cdot \prod_{z \in N_{\mathbb{A}}(x)\setminus\{u,y\}} \bar{q}_\lambda^z \le \pi_\ell \frac{\gamma(x)}{\Gamma_{\mathbb{A}}^{\circ}(u)} \cdot \frac{\Gamma_{\mathbb{A}}^{\circ}(x)}{\Gamma_{\mathbb{A}}^{\circ}(u)} \prod_{z \in N_{\mathbb{A}}(x)\setminus\{u,y\}} \bar{q}_\lambda^z \tag{21}$$

Now, using Lemma A.1, observe that $\sum_{x \in N^3(u)} \frac{\gamma(x)}{\Gamma(x)} \le \alpha^3$ and therefore:

$$\sum_{x \in X_m^{u,v}} \gamma(x) = \sum_{x \in X_m^{u,v}} \frac{\gamma(x)}{\Gamma(x)} \Gamma(x) \le \sum_{x \in X_m^{u,v}} \frac{\gamma(x)}{\Gamma(x)} \left(\Gamma_{\mathbb{A}}^{\circ}(x) + \frac{1}{2}\right) \le 10(m+1)\alpha^3\Gamma_{\mathbb{A}}^{\circ}(u) + \alpha^3 \tag{22}$$

We first look at the case $m = 0$.

Case 1a: $\Gamma_{\mathbb{A}}^{\circ}(u) > 2$, $m = 0$. Then, by (22), $\sum_{x \in X_m^{u,v}} \gamma(x) \le 11\alpha^3\Gamma_{\mathbb{A}}^{\circ}(u)$. Also, (20) can be upper bounded by $\frac{4}{3}\pi_\ell \frac{\gamma(x)}{\Gamma_{\mathbb{A}}^{\circ}(u)}$ (since $\bar{q}_\lambda^y \le 1$). For (21) we note that $\Gamma_{\mathbb{A}}^{\circ}(x) \le 10\Gamma_{\mathbb{A}}^{\circ}(u)$ and get combined that $\mathbf{P}(E_\lambda^x|B_\lambda^{u,v}) \le 10\pi_\ell \frac{\gamma(x)}{\Gamma_{\mathbb{A}}^{\circ}(u)}$, regardless whether $x \in N_{\mathbb{A}}^1(v)$ or not.

$$\mathbf{P}(F_{\lambda,0}^{u,v}|B_\lambda^{u,v}) \ge \prod_{x \in X_0^{u,v}} \left(1 - 10\pi_\ell \frac{\gamma(x)}{\Gamma_{\mathbb{A}}^{\circ}(u)}\right) \ge e^{-2\frac{10\pi_\ell}{\Gamma_{\mathbb{A}}^{\circ}(x)} \sum_{x \in X_0^{u,v}} \gamma(x)} \overset{(22)}{\ge} e^{-220\alpha^3\pi_\ell} \overset{\pi_\ell \le \frac{1}{2}}{=} \Omega(1).$$

Case 1b: $\Gamma_{\mathbb{A}}^{\circ}(u) > 2$, $m \ge 1$. We show that, w.c.p., all $x \in \bigcup_{m \ge 1} X_m^{u,v}$, that listen on channel $\mathcal{A}_\lambda$, have a collision. For E_λ^x to happen, exactly one of x's neighbors has to broadcast. We look again first at $x \in N_{\mathbb{A}}(v) \setminus \{u\}$. We plug in the values for $\bar{q}_\lambda^y$ into (20) and we use that $\Gamma_{\mathbb{A}}^{\circ}(x) \ge 10m\Gamma_{\mathbb{A}}^{\circ}(u)$ to get

$$\mathbf{P}(E_\lambda^x|B_\lambda^{u,v}) \le \frac{4}{3}\pi_\ell \frac{\gamma(x)}{\Gamma_{\mathbb{A}}^{\circ}(u)} \prod_{y \in N_{\mathbb{A}}(x)\setminus\{u,v\}} \left(1 - (1-\pi_\ell)\frac{\gamma(y)}{\Gamma_{\mathbb{A}}^{\circ}(u)}\right) \le \frac{64}{27}\pi_\ell \frac{\gamma(x)}{\Gamma_{\mathbb{A}}^{\circ}(u)} e^{-\frac{1}{2}\frac{\Gamma_{\mathbb{A}}^{\circ}(x)}{\Gamma_{\mathbb{A}}^{\circ}(u)}} \le 3\pi_\ell \frac{\gamma(x)}{\Gamma_{\mathbb{A}}^{\circ}(u)} e^{-5m}$$

For $x \in N^3_{\mathbb{A}}(v) \setminus N^1_{\mathbb{A}}(v)$ we do the same with (21). We also use $\Gamma^\circ_{\mathbb{A}}(x) \le 10(m+1)\Gamma^\circ_{\mathbb{A}}(u)$ for our calculations:

$$\begin{aligned}\mathbf{P}(E^x_\lambda | B^{u,v}_\lambda) &\le \pi_\ell \frac{\gamma(x)}{\Gamma^\circ_{\mathbb{A}}(u)} \cdot \frac{\Gamma^\circ_{\mathbb{A}}(x)}{\Gamma^\circ_{\mathbb{A}}(u)} \prod_{z \in N_{\mathbb{A}}(x)\setminus\{u,y\}} \left(1 - (1-\pi_\ell)\frac{\gamma(z)}{\Gamma^\circ_{\mathbb{A}}(u)}\right) \\ &\le \frac{16}{9} 10(m+1)\pi_\ell \frac{\gamma(x)}{\Gamma^\circ_{\mathbb{A}}(u)} e^{-\frac{1}{2}\frac{\Gamma^\circ_{\mathbb{A}}(x)}{\Gamma^\circ_{\mathbb{A}}(u)}} \le 40 m \pi_\ell \frac{\gamma(x)}{\Gamma^\circ_{\mathbb{A}}(u)} e^{-5m}\end{aligned}$$

For all $x \in X^{u,v}_m$ we therefore have $\mathbf{P}(E^x_\lambda|B^{u,v}_\lambda) \le 40 m \pi_\ell \frac{\gamma(x)}{\Gamma^\circ_{\mathbb{A}}(u)} e^{-5m}$. Recall that by (22) we can upper bound $\sum_{x \in X^{u,v}_m} \gamma(x)$ by $22 m \alpha^3 \Gamma^\circ_{\mathbb{A}}(u)$.

For $F^{u,v}_{\lambda, m\ge 1} := \bigcap_{m \ge 1} F^{u,v}_{\lambda,m}$ we get

$$\begin{aligned}\mathbf{P}(F^{u,v}_{\lambda,m\ge1} | B^{u,v}_\lambda) &\ge \prod_{m\ge1} \prod_{x \in X^{u,v}_m} (1 - \mathbf{P}(E^x_\lambda | B^{u,v}_\lambda)) \ge \prod_{m\ge1} \prod_{x\in X^{u,v}_m} \left(1 - 40 m \pi_\ell \frac{\gamma(x)}{\Gamma^\circ_{\mathbb{A}}(u)} e^{-5m}\right) \\ &\ge \prod_{m\ge1} e^{-80 m \pi_\ell e^{-5m} \sum_{x\in X^{u,v}_m} \frac{\gamma(x)}{\Gamma^\circ_{\mathbb{A}}(u)}} \ge \prod_{m\ge1} e^{-1760 \pi_\ell \alpha^3 m^2 e^{-5m}} \\ &\ge e^{-1760\pi_\ell \alpha^3 \sum_{m\ge1} m^2 e^{-5m}} \ge e^{-12\pi_\ell\alpha^3} = \Omega(1).\end{aligned}$$

Let us now define $F^{u,v}_\lambda$ as $F^{u,v}_{\lambda,m=0} \cap F^{u,v}_{\lambda,m\ge1}$, then we get:

$$\mathbf{P}(D^{u,v}_\lambda) = \mathbf{P}(B^{u,v}_\lambda \cap F^{u,v}_\lambda) = \mathbf{P}(B^{u,v}_\lambda)\mathbf{P}(F^{u,v}_{\lambda,m=0}|B^{u,v}_\lambda)\mathbf{P}(F^{u,v}_{\lambda,m\ge1}|B^{u,v}_\lambda) = \Omega(\mathbf{P}(B^{u,v}_\lambda))$$

With the results from above, this concludes the analysis for the case of $\Gamma^\circ_{\mathbb{A}}(u) > 2$.

Case 2: $\Gamma^\circ_{\mathbb{A}}(u) \le 2, \lambda = 1$ and $2^{-\lambda} = 1/2$.

For this case we redefine $X^{u,v}_m$.

$$X^{u,v}_m := \left\{x \in N^3_{\mathbb{A}}(v) \setminus \{u,v\} : 10m \le \Gamma^\circ_{\mathbb{A}}(x) < 10(m+1)\right\}$$

Analogously to (20) and (21) we get (with $2^{-\lambda} = 1/2$):

$$x \in N_{\mathbb{A}}(v)\setminus\{u\}: \qquad \mathbf{P}(E^x_\lambda|B^{u,v}_\lambda) = \mathbf{P}(E^x_1|B^{u,v}_1) \le \frac{2}{3}\pi_\ell\gamma(x) \prod_{y\in N_{\mathbb{A}}(x)\setminus\{u,v\}} \bar{q}^y_\lambda \tag{23}$$

$$x \notin N^1_{\mathbb{A}}(v): \qquad \mathbf{P}(E^x_\lambda|B^{u,v}_\lambda) = \mathbf{P}(E^x_1|B^{u,v}_1) \le \frac{\pi_\ell}{4}\gamma(x)\Gamma^\circ_{\mathbb{A}}(x) \prod_{z\in N_{\mathbb{A}}(x)\setminus\{u,y\}} \bar{q}^z_\lambda \tag{24}$$

Like with (22), Lemma A.1 gives us that

$$\sum_{x\in X^{u,v}_m} \gamma(x) \le 10(m+1)\alpha^3 + \frac{1}{2}\alpha^3 \le 11(m+1)\alpha^3. \tag{25}$$

Case 2a: $\Gamma^\circ_{\mathbb{A}}(u) \le 2, m = 0$.

From 25 we get $\sum_{x\in X^{u,v}_0} \gamma(x) \le 11\alpha^3$. With $\bar{q}^y_\lambda = \bar{q}^y_1 \le 1$ and $\Gamma^\circ_{\mathbb{A}}(x) \le 10$ we get for all $x \in X^{u,v}_0$ that $\mathbf{P}(E^x_\lambda|B^{u,v}_\lambda) = \mathbf{P}(E^x_1|B^{u,v}_1) \le 3\pi_\ell\gamma(x)$.

Thus,

$$\mathbf{P}(F^{u,v}_{\lambda,0}|B^{u,v}_\lambda) \ge \prod_{x\in X^{u,v}_0} (1 - 3\pi_\ell\gamma(x)) \ge e^{-66\pi_\ell\alpha^3} = \Omega(1)$$

Case 2b: $\Gamma^\circ_{\mathbb{A}}(u) \le 2, m \ge 1$.

Note that since $\lambda = 1$ we have $q_1^w = (1-\pi_\ell)\gamma(w)2^{-\lambda} \ge \gamma(w)/4$ and therefore $\bar{q}_1^w \le e^{-\gamma(w)/4}$. We use (23) and (24) to show that for *any* $x \in X_m^{u,v}$

$$\mathbf{P}(E_\lambda^x | B_\lambda^{u,v}) \le \frac{1}{4}\pi_\ell \gamma(x) 20m \left(\frac{4}{3}\right)^2 \prod_{z \in N_{\mathbb{A}}(x)} \bar{q}_\lambda^z \le 10m\pi_\ell\gamma(x)e^{-2.5m}$$

As indicated by (25), $\sum_{x \in X_m^{u,v}} \gamma(x) \le 22m\alpha^3$ for $m \ge 1$.
For $F_{\lambda,m\ge1}^{u,v} = \bigcap_{m\ge1} F_{1,m}^{u,v}$ we thus get again

$$\begin{aligned}
\mathbf{P}(F_{1,m\ge1}^{u,v}|B_1^{u,v}) &\ge \prod_{m\ge1}\prod_{x\in X_m^{u,v}}(1-\mathbf{P}(E_1^x|B_1^{u,v})) \ge \prod_{m\ge1}\prod_{x\in X_m^{u,v}}\left(1-10m\pi_\ell\gamma(x)e^{-2.5m}\right)\\
&\ge \prod_{m\ge1} e^{-20m\pi_\ell e^{-2.5m}\sum_{x\in X_m^{u,v}}\gamma(x)} \ge \prod_{m\ge1} e^{-440m\pi_\ell\alpha^3 e^{-5m}} \ge e^{-440\pi_\ell\alpha^3\sum_{m\ge1} me^{-2.5m}}\\
&\ge e^{-50\pi_\ell\alpha^3} = \Omega(1).
\end{aligned}$$

Analogously to the case $\Gamma^\circ_{\mathbb{A}}(x) > 2$ it holds that $\mathbf{P}(D_1^{u,v}) = \Omega(\mathbf{P}(B_1^{u,v}))$.

Let c be the constant such that $\mathbf{P}(D_\lambda^{u,v}) \ge c\mathbf{P}(B_\lambda^{u,v})$ for any $v \in N_{\mathbb{A}}(u)$. Then, since $D_\lambda^u = \biguplus_{v\in N_{\mathbb{A}}(v)} D_\lambda^{u,v}$ and $B_\lambda^u = \biguplus_{v\in N_{\mathbb{A}}(v)} B_\lambda^{u,v}$, it most hold that $\mathbf{P}(D_\lambda^u) \ge c\mathbf{P}(B_\lambda^u)$. Also, $\mathbf{P}(D^u) \ge \mathbf{P}(D_\lambda^u)$ and since by (18) $\mathbf{P}(B_\lambda^u) = \Omega(\mathbf{P}(B^u))$, we get that $\mathbf{P}(D^u) = \Omega(\mathbf{P}(B^u))$.

This finishes the proof for the second part of the claim, equations (11). □

Corollary 4.9. *Let $B(r)^{u,v}$ be the event that in round r node $u \in \mathbb{A}_r$ receives a message from one of its neighbors $v \in \mathbb{A}_r$, neither of them neighboring any leader, herald or herald candidate in the 5th or 6th round of its handshake protocol. Let $\hat{H}(r')^{u',v'}$ be the event that at the beginning of round r' node $u' \in \mathbb{H}_{r'}$ and $v' \in \mathbb{L}_{r'}$ form a good leader-herald pair and that $\mathbb{H}' \cap N^3(u') = \emptyset$, i.e., there are no herald candidates in the 3-neighborhood of u'. Then*

$$\mathbf{P}(\hat{H}(r+8)^{u,v}|B(r)^{u,v}) = \Omega(1) \tag{26}$$

Proof. According to Lemma 4.8, since $D^u \subset B^u$, it holds that $\mathbf{P}(D^u|B^u) = \Omega(1)$. Thus, if $B(r)^{u,v}$ happens, where u receives the message on channel $\mathcal{A}_i$, then, w.c.p., no other node $w' \in N^3(u)\setminus\{u\}$ (which includes v's neighborhood) receives a message on $\mathcal{A}_i$. If no other node $w' \in N^3(u)\setminus\{u\}$ receives a message on $\mathcal{A}_i$, that means that all such nodes w' either not operate on $\mathcal{A}_i$, or, if they do, none or at least 2 of their neighbors $w'' \in N^4(u)$ send a message on $\mathcal{A}_i$. Using Lemma 4.6.(4) by setting $k = 4$, π_ℓ appropriately small and conditioning on the set $S = S^l \cup S^b \cup S^n$ of nodes as listening/broadcasting on channel $\mathcal{A}_i$ or not operating on $\mathcal{A}_i$ at all (denoted by event $Z := \mathcal{S}_{\neg i}^n \wedge \mathcal{S}_i^b \wedge \mathcal{S}_i^l$), we get that $\mathbf{P}(H_{\neg i}|D^u) = \mathbf{P}(H_{\neg i}|Z) > 1/2$.

If in round r both events $H_{\neg i}$ and D^u happen (where u gets its message on channel $\mathcal{A}_i$), then in round $r+1$ nodes u and v meet on channel $\mathcal{H}$ and can perform the first round of the handshake protocol without any interruption from any other nearby nodes. By applying Lemma 4.6.(1) multiple times, in rounds $r+1$ to $r+7$, w.c.p., no other new herald is created in $N^3(u)$. Thus, w.c.p., by the beginning of round $r+8$, u and v have emerged from the handshake as a good leader-herald pair. □

We next recall the analysis done for the so-called *Handshake* and *Red-Blue Game*, which shows that the creation of good pairs implies progress. All lemmas for the handshake protocol and the red-blue protocol have been proven in [9] and their proofs are almost unaffected by the changes made to the algorithm.

4.5 Handshake Protocol—Nodes in States $\mathbb{L}'$ and $\mathbb{H}'$

Lemma 4.10. *In round r consider two leader-herald pairs (l_1, h_1) and (l_2, h_2) and suppose that the pairs started their most recent handshakes in rounds r_1 and r_2, $r_1 \leq r_2$, respectively. Say that edge e is* crossing *if one of its endpoints is in $\{l_1, h_1\}$ and its other endpoint is in $\{l_2, h_2\}$. Then, either no crossing edge exists or exactly one of the following conditions is satisfied: (1) $r_1 = r_2$ and crossing edges are $\{l_1, l_2\}$ and/or $\{h_1, h_2\}$, (2) $r_2 = r_1 + 2$ and the only crossing edge is $\{l_1, h_2\}$.*

Lemma 4.10 corresponds to Lemma 8.10 in [9] and in the proof there it was made use of the fact that in the original red-blue game in every second round both leader and herald blocked channel $\mathcal{H}$. But the new red-blue game features exactly the same mechanic. Everything else in the proof stays the same as the handshake protocol has not changed at all, hence we omit the proof here.

4.6 Red-Blue Protocol—Nodes in States $\mathbb{L}$ and $\mathbb{H}$

Lemma 4.11. *If a pair (l, h) is good in round r and they started their first red-blue game in round r', then by the end of round $r' + \tau_{\text{red-blue}} = r' + \Theta(\log n)$, w.h.p., either*

- *the related leader l joins the MIS, or*
- *a node $v \in N(l) \cup N(h)$ joins the MIS by increasing its lonely counter above τ_{lonely}.*

This lemma and its proof are exactly the same as in [9], hence we omit any proof. Note that the lemma stresses that a good leader does not necessarily become an MIS node. To see how the second case can happen, assume that some node v is in the herald filter already for a long time, with its lonely counter almost reaching the threshold. If now a neighboring node in G makes it out of the decay filter, and, after increasing its activity value for a while, communicates successfully with another node. This can happen before v notices the existence of u, and shortly after u becomes a leader, v could join the MIS via the lonely counter.

Leaders of bad pairs do not have a justified claim for being MIS nodes, so we do not want them to hinder progress. But they do that by decreasing activity values of their neighbors. The next lemma shows that bad pairs do not last for very long, thus not causing a long stagnation.

Lemma 4.12. *Consider a node v and suppose that in an arbitrary round r, there is a leader or herald of a bad pair in $N^3(v)$. Then, w.c.p., in round $r + 16$, no node in $N^3(v)$ is in state $\mathbb{H}'$ and all leaders and heralds are part of a good pair.*

The proof for Lemma 4.12 remains the same, too, except that in the original version round $r + 12$ was stated. This change is due to the fact that the length of a single red-blue game did increase from 6 to 8. It does not affect the proving method and we again omit any proof.

4.7 Joining the MIS—Nodes in States $\mathbb{M}$ and $\mathbb{E}$

Property 1 (P). *The set $\mathbb{M}$ is an independent set at all times.*

This intuitive assumption is needed for some of the upcoming statements; it is clearly true at the beginning of the algorithm, when $\mathbb{M} = \emptyset$. We show in Lemma 4.18, that if (P) is violated, then w.h.p. a contradiction occurs. The next lemma makes sure that nodes in $N(\mathbb{M})$ soon learn of their coverage.

Lemma 4.13. *Assume (P) holds. Let v be a node that enters state $\mathbb{M}$ at time t. Let w be a node in $N_G(v)$ that is awake at time $t' \geq t$ and, if $w \in \mathbb{L} \cup \mathbb{H}$, that it is at most in round $\frac{9}{10}\tau_{\text{red-blue}}$ of its corresponding red-blue game. Then by time $t' + \tau_{\text{notification}} = t' + O(\log n)$, w.h.p., w is in state $\mathbb{E}$.*[10]

[10] Note that this lemma also considers nodes w from the decay filter.

Proof. We proof the statement by showing that within $O(\log n)$ rounds a non-sleeping node w receives a message from $v \in \mathbb{M}$ on some channel $\mathcal{R}_i$. Since nodes in the decay filter listen on the report channels in any given round with probability at least $1/2$, and therefore at least half as often as nodes in the herald filter, we restrict our analysis to nodes in the herald filter and proof the statement for a time bound of $\tau_{notification}/2$. State changes from the decay filter to the herald filter do not affect the analysis because of the very same reason. Also note that w is *unable* to move to state $\mathbb{L}$ nor $\mathbb{H}$ after time $t+4$, since v disrupts any handshake in its neighborhood by sending at least once on $\mathcal{H}$ in every set of two consecutive rounds.

Case w $\in \mathbb{A}$. Consider some round $t'' \geq t'$. If w is in state $\mathbb{A}$, w.c.p., it also is in that state in round $t''+1$. Further, w.c.p., MIS node v has its variable *enforce* set to false in round $t''+1$, and thus, w.c.p., broadcasts on some channel $\mathcal{R}_i$ in that round. Assume w does not neighbor any bad herald or bad leader, then it neighbors at most α good leaders, α^2 good heralds, and α MIS nodes. To see this, note that while good heralds are allowed to be adjacent to each other, each has a neighboring good leader, and thus the number of adjacent good heralds in the direct neighborhood of w is upper bounded by the number of good leaders in $N^2(w)$, and thus by α^2. We can therefore upper bound the number of adjacent good leaders, good heralds, or MIS nodes by $3\alpha^2$. The probability that v chooses channel $\mathcal{R}_i$ while no good herald, good leader or other MIS node neighboring w chooses to operate on the same channel $\mathcal{R}_i$, is constant, since $n_{\mathcal{R}}$ is greater than $3\alpha^2$, but still a constant. With probability at least $\frac{1}{2n_{\mathcal{R}}} = \Omega(1)$, w listens on $\mathcal{R}_i$ in round $t''+1$, and therefore, w learns of v's state with constant probability.

Now let there be bad pairs in w's neighborhood in round t'' and w be in $\mathbb{A}$. Then, w.c.p., by Lemma 4.12, 16 rounds later w is still/again in state $\mathbb{A}$ while all bad pairs in $N^2(w)$ are knocked out and no new bad pair has been created due to Lemma 4.6.(1). As before, w learns of v's state w.c.p. after $O(1)$ rounds.

Let now w be in different states. If w is in $\mathbb{H}'$ (but at most in the fourth round of the handshake protocol) or in $\mathbb{L}'$, its handshake will fail due to v's routine of disrupting channel $\mathcal{H}$ at least once every 2 rounds, reverting w back to state $\mathbb{A}$. For the cases of w being either in $\mathbb{L}$, $\mathbb{H}$ or in the last two rounds of the handshake protocol as a herald candidate ($\mathbb{H}'$) (which we denote by $\mathbb{H}'_{5,6}$ in this proof), note that if w ever moves to state $\mathbb{A}$ from these states, it is unable to return. If w is in $\mathbb{L}$, in each red-blue game, there is a constant probability of $\Theta(1)$ that v disrupts the game, bringing w back to state $\mathbb{A}$. Instead, if w is in $\mathbb{H}$, during each red-blue game, there is a $\Omega(1/n_{\mathcal{R}})$ probability that v operates on the same channel $\mathcal{R}_{meet}$ as w does in round 8 of its respective red-blue game, disrupting the ongoing red-blue game and sending w back to state $\mathbb{A}$. Due to our condition of w not being too far in its red-blue protocol, we can choose $\tau_{red\text{-}blue}$ large enough to make sure that w, whether it is in state $\mathbb{L}$ or $\mathbb{H}$, hears from v w.h.p. before it can join the MIS itself. If w is in the last two rounds of the handshake protocol, then this case can be reduced to w being a herald.

Thus, if w is in $\mathbb{L} \cup \mathbb{H} \cup \mathbb{H}'_{5,6}$, then it leaves this set within $O(\log n)$ rounds and never returns to it. Thus let us assume that w is not in this set (anymore).

Clearly, if w is not in $\mathbb{A}$ in any round $t'' \geq t'$, w.c.p., it returns to $\mathbb{A}$ in $O(1)$ rounds. And as argued above, if bad pairs are in w's neighborhood, they get eliminated w.c.p. in $O(1)$ rounds as well. Choosing $\tau_{notification} = O(\log n)$ sufficiently large and applying a Chernoff bound proves the statement. ☐

4.8 Progress and Runtime

In Lemma 8.13 of [9] we have shown that once a good leader-herald pair is created, its leader (or another node in distance 2) joins the MIS within $O(\log n)$ rounds. Also, we used the fact that within close proximity of *fat* nodes (which exist in any δ-neighborhood of any node in the herald filter) such solitary pairs are created with a constant probability. In the algorithm in [9] it might happen that after a good pair is created within radius δ of some node u, the only fat node within $N^{\delta}(u)$ is close to that pair. A good pair blocks the creation of other pairs around it, so progress might be stalled until the leader of the good pair joins the MIS, causing it to eliminate its neighbors (and therefore their activity) and finally, forcing the local condition of fatness to move to a different area of the graph.

Here we changed the algorithm to take care of this potential stagnancy issue. We want the attribute of fatness to move away from a good pair long before the leader joins the MIS. More precisely, a node not neighboring good pairs should become fat within $o(\log n)$ rounds. For this we require good pairs to reduce the activity levels of their neighborhoods. However, a leader-herald pair does not know whether it forms a good pair or a bad one before the $\tau_{\textit{red-blue}} = \Theta(\log n)$ red-blue games are over. The idea to deal with this difficulty is the following. Good pairs manage in expectation within $O(\log\log n)$ rounds to reduce their neighborhood's activity far enough such that most of the time those nodes can be considered inactive. Bad pairs, however, last for only a constant number of rounds in expectation, and are created rarely enough[11] for affected nodes to recover their lost activity quickly. In other words, the longer a node is a leader, the more likely it is that this node is good.

Careful analysis allows us to transform these observations into high probability results.

In the following $\gamma(u,t)$ denotes the activity level of node u in round t. Also, let ε be a constant smaller than 1—about 0.1 is sufficiently small for the upcoming proofs.

Lemma 4.14. *Let t be a time at which a node $u \notin N^1(\mathbb{M})$ is in the herald filter. Then, w.h.p., one of following holds:*
(a) Within $\tau_{\textit{progress}} = O(\log^2 n/\mathcal{F} + \log n)$ rounds, $u \in N^1(\mathbb{M})$, or
(b) $|\{t' \in [t+1, t+\tau_{\textit{progress}}] : \gamma(u,t') = 1/2\}| \geq (1-\varepsilon)\tau_{\textit{progress}}$.

Proof. Initially assume that the lemma allows in addition to conditions (a) and (b) the following:
(c) or within $\tau_{\textit{progress}} = O(\log^2 n/\mathcal{F} + \log n)$ rounds, there is a good leader in $N^2(u)$.
We prove the statement by contradiction, thus assume that neither (a), (b) nor (c) holds.

Let $t < T_1 < T_2 < \cdots < T_m \leq t + \tau_{\textit{progress}} - \tau_{\textit{red-blue}}$ be the rounds in which a respective series of leaders or heralds $v_1, v_2, \ldots, v_m \in N(u)$ neighbors u for the first time, i.e., v_i successfully finished its handshake in round $T_i - 1$ and managed to reach/was reached by a single node $v'_i \in N^2(u)$ in round $T_i - 7$ on some channel $\mathcal{A}_{k_i}$. For $T_i < T_j$ it can hold that $v_i = v_j$ — in this case node v_i moved from $\mathbb{L} \cup \mathbb{H}$ to state $\mathbb{A}$ in the time interval $[T_i, T_j]$. Also, at time T_i more than one leader/herald could neighbor u for the first time, in which case we let v_i be any of these. For the pair (v_i, v'_i) denote with l_i the leader and with h_i the corresponding herald.

Assume that all corresponding pairs (l_i, h_i) are bad pairs as otherwise the lemma would be trivially fulfilled.

Fix i. l_i can become a bad leader in round T_i, if in round $T_i - 7$ a node $l' \in N_{\mathbb{A}}(l_i)$ reaches another node $h' \in \mathbb{A}$ in that round and manages to get through its handshake protocol as well. For (l_i, h_i) to get through their handshake protocol, they cannot neighbor a leader or a herald or a herald candidate in round 5 or 6 at that time. By Corollary 4.9, the probability for (l_i, h_i) to turn out a bad pair is in $1 - \Omega(1)$. Another way for l_i to become a bad leader is if in round $T_i - 5$ a node $h' \in N(l_i)$ successfully receives a message. By Lemma 4.6.(1) this happens only with probability $O(\pi_\ell)$. Hence, in both cases, the probability for l_i to be a bad leader conditioned on the event that l_i becomes a leader in round T_i, is at most $1 - \Omega(1)$. Since we assume that all leaders $l_1, l_2, \ldots, l_m$ are bad leaders, a Chernoff bound then gives us that, w.h.p., $m = O(\log n)$. Assume this is the case. Let X_i be the random variable that measures the number of red-blue games leader l_i survives before it either becomes part of a good pair or it gets knocked out. If l_i becomes part of a good pair then the extended lemma statement would be fulfilled, hence we assume otherwise. Note that, w.h.p., l_i (like any other node) cannot finish its red-blue protocol as a bad leader and join the MIS.

So, for all i, let us assume that l_i gets knocked out of state $\mathbb{L}$ while being part of a bad pair. According to Lemma 4.12 this happens in each red-blue game with constant probability, i.e., $\mathbb{E}[X_i] = \Theta(1)$. Let Y be the number of rounds in which at least one node in $N(u)$ plays the red-blue game. Then, since $m = O(\log n)$, applying Chernoff once more we have for $Y \leq X := \sum_{i=1}^{m} X_i$ that $\mathbb{E}[Y] \leq \mathbb{E}[X] = O(\log n)$. If u

[11] controlled by reducing the parameter π_ℓ

always picks a channel R_{meet} for the corresponding rounds on which a neighboring bad pair communicates, then u reduces its activity by the factor $\sigma_\ominus$ in each such round. This totals to a reduction of at most $\sigma_\ominus^Y$, spread over $\tau_{progress}$ rounds. But since $Y = O(\log n)$ it also takes at most $Y \cdot \log_{\sigma_\oplus} \sigma_\ominus = Y \cdot 20\bar{m} = \Theta(Y) = O(\log n)$ rounds to recover from those activity reductions. Choosing parameter $\tau_{progress}$ large enough, $|\{t' \in [t+1, t+\tau_{progress}] : \gamma(u,t') = 1/2\}| \geq (1-\varepsilon)\tau_{progress}$ holds with high probability.

What remains is to remove condition (c). But this follows from the simple fact that if a good leader is created in $N^2(u)$, then within $\tau_{red\text{-}blue}$ rounds it either joins the MIS or it gets knocked out by an MIS node due to Lemma 4.11. The latter can happen at most α^3 times. Thus we can omit condition (c) by extending $\tau_{progress}$ additively by $\tau_{red\text{-}blue}$ and then by a factor of α^3. □

Next we upper bound the number of rounds in which *any* neighbors of good pairs within distance δ from u manage to exceed the activity threshold γ_{low}.

Definition 4.15. *For a node u and a round r let $I(u,r)$ be the event that*

- all *nodes $x \in N^\delta(u)$, which neighbor an MIS node, are in state $\mathbb{E}$, and*
- all *nodes $x \in N^\delta(u)$, which neighbor a good herald h or good leader l, $h, l \in (\mathbb{H} \cup \mathbb{L}) \setminus N(\mathbb{M})$, have $\gamma(x) \leq \gamma_{low} = \sqrt{\gamma_{\min}} = \log^{-12} n$ and are neither bad leaders nor bad heralds.*

Lemma 4.16. *Assume that (P) holds. Further, let $\bar{r}$ be a round in which node u is in the herald filter and set $J := [\bar{r}+1, \bar{r}+\tau_{progress}]$. Then, w.h.p., one of the following holds:*

- *Within $\tau_{progress} = O(\log n)$ rounds, there is an MIS node in $N^1(u)$, or*
- $|\{r \in J : I(u,r) \textit{holds}\}| \geq (1-\varepsilon)\tau_{progress}$.

Proof. Initially, $\tau_{progress}$ is chosen large enough to comply with Lemma 4.14.

The proof is divided into four parts. In the first part we show that there are only $\operatorname{polylog} n$ many nodes for which $I(u,r)$ can be violated—i.e., nodes that neighbor good leaders, good heralds or MIS nodes. In the second part we show that each such node hears from one of these neighbors every $O(1)$ rounds with constant probability. Since neighbors of MIS nodes are immediately eliminated upon hearing a message from them, the case of a node neighboring a good pair is the more difficult one. In the third part we argue that a node neighboring good pairs reduces its activity value to $\gamma_{\min}$ at least once within $O(\log\log n)$ rounds with considerable probability: $1 - \log^{-c} n$ for some constant $c > 1$. In the last part we combine those results to show that within $\tau_{progress}$ rounds, w.h.p., $I(u,r)$ is only violated for a small constant fraction of those rounds.

First part. To count the number of nodes that can violate $I(u,r)$, we have to count the number of good heralds, good leaders and good MIS nodes in $N^{\delta+1}(u)$ in round r; we denote that latter set by $W_r \subseteq (\mathbb{L} \cup \mathbb{H} \cup \mathbb{M}) \cap N^{\delta+1}(u)$ and their neighboring nodes in the herald filter by $N_r := N(W_r)$. Due to (P) and Lemma 4.11, w.h.p. at all times all MIS nodes form an independent set, as do all leaders of good pairs, so there are no more than $2\alpha^{\delta+1}$ of these in W_r. The number of heralds of good pairs in $N^{\delta+1}(u)$ is at most the number of good leaders in $N^{\delta+2}(u)$, so in total $|W_r|$ amounts to at most $3\alpha^{\delta+2} \leq 3\alpha^2\sqrt{\log n} = o(\log n)$. Due to the guarantee we get from the decay filter, each such node has at most $O(\log^3 n)$ neighbors in the herald filter, so $N_r < \log^4 n$. Note that even over $\tau_{progress}$ rounds the total amount of those nodes cannot exceed $o(\log^6 n)$.

Second part. An MIS node has to reach all its neighbors only once each, in order for them to fulfill the requirement for $I(u,r)$. Leaders and heralds on the other hand need to inform their neighbors multiple times and continuously over the course of $\tau_{progress}$ rounds. In most cases a node $v \in N_r$ has a constant probability to hear of at least one of its neighbors in W_r within a $O(1)$-round time interval. More precisely, we claim that within $O(\log\log n)$ rounds, in expectation, v reduces its activity by a factor polylogarithmic in n. We prove this by analyzing the 3 types of neighbors v can have in W_r and the states v can be in.

$v \in \mathbb{A}$: At first let v neighbor an MIS node s. s might be forced to broadcast on $\mathcal{H}$ in round r, but then with probability $1/4$ it sends on one of the report channels $\mathcal{R}_k$ in round $r+1$. Also with probability at least $1/2$, v does not act on one of the active channels $\mathcal{A}_i$ in round r and therefore either gets knocked out or is also in state $\mathbb{A}$ in round $r+1$, where it listens on some report channel with probability at least $1/2$. Thus, with probability at least $1/16$, in round $r' \in [r, r+1]$, s sends on a report channel and v listens on a report channel. If that is the case, since $|W_{r'} \cap N^1(v)| \leq n_{\mathcal{R}}$, with probability at least $\frac{1}{en_{\mathcal{R}}}$, v and s act on the same channel $\mathcal{R}_k$ while no other neighbor of v in $W_{r'}$ operates on $\mathcal{R}_k$. Once that happens, v moves to state $\mathbb{E}$.

Thus let us now assume that v does not neighbor an MIS node, but at least one good leader l. With similar reasoning, there is a round $r' \in [r, r+7]$ in which l broadcasts on some channel $\mathcal{R}_k$ as a apart of its red-blue game, unless l transitioned to state $\mathbb{M}$ or $\mathbb{E}$—the first case we already covered, the second does not happen with probability at least $1/2$. With probability at least 2^{-7}, v is still or again in state $\mathbb{A}$ in that round r'. With similar argumentation as above, v then gets l's message with probability at least $\frac{1}{2en_{\mathcal{R}}}$. In total, in each 8-round interval, with probability at least $\frac{1}{2^8 en_{\mathcal{R}}}$, v receives a message from l.

If v does not neighbor a good leader nor an MIS node, but at least one good herald h, then the same logic applies for an 8-round interval, unless h gets knocked out—which only happens if its leader joins the MIS or another MIS node is created nearby. But, with probability at least $1/2$, h nor its leader gets knocked out by a neighboring MIS node. If its leader joined the MIS, then just a new MIS node has been created in $N^2(v)$, which can happen at most α^2 times and therefore delay v to hear from any of its neighbors by at most $O(\alpha^2)$ rounds during the whole execution of the algorithm. For simplicity we ignore those $O(1)$ rounds. As above, with probability at least $\frac{1}{2^8 en_{\mathcal{R}}}$, v hears from h in any 8-round interval.

In total we get that for every $2^{11} en_{\mathcal{R}}$ rounds that v spends in $\mathbb{A} \cap N_r$, in expectation it hears at least once from one of its neighbors in W_r. At this point we fix $\bar{m}$ to be eight times as large, i.e., $2^{14} en_{\mathcal{R}}$. Note that $\bar{m}$ is a constant depending on α only, via $n_{\mathcal{R}}$.

$v \in \mathbb{H} \cup \mathbb{L}$: As long as v is in one of these states, v hears from the node it partnered up with every 8 rounds as long as the pair remains. In the algorithm it is accounted for that by decreasing $\gamma(v)$ every round by a factor of $\sigma_{\oplus}^{20}$.

$v \in \mathbb{L}' \cup \mathbb{H}'$: If v was already in $\mathbb{L}'$ or $\mathbb{H}'$ when it joined N_r, then it might manage to finish its handshake, in which case we refer to the previous case. Otherwise, if its handshake gets disrupted, then v returns to state $\mathbb{A}$ at latest 6 rounds later, which we also already covered. This also implies, that a node switching between states $\mathbb{A}$ and $\mathbb{H}'$ or $\mathbb{L}'$ spends at least one eight of these rounds in $\mathbb{A}$.

We get that v hears from one of its neighbors in W_r at least once every $\bar{m}$ rounds, irrelevant of its own state. Recalling that $\sigma_{\ominus} = \sigma_{\oplus}^{20\bar{m}}$, v reduces its activity each 8 rounds by a factor of at least $\sigma_{\oplus}^{152}$ in expectation.

Third part. Let r_i be the round in which the ith MIS node or good leader-herald pair in $N^{\delta}(u)$ is created and we talk of event $\mathcal{E}_{r_i}$; if more than one is created in a single round, we ignore additional ones. In round r_i all nodes $v \in N_{r_i}$ have $\gamma(v) \leq 1/2$. Let us assume that no other good pair or MIS node is created for $\Omega(\log \log n)$ rounds, as otherwise we assume all rounds between r_i and r_{i+1} as *violated*—more details on that in part four. We want to show that all nodes $v \in N_{r_i}$ decrease their activity quickly. For as long as v is part of a pair, this reduction is guaranteed by design of the algorithm, so we also only consider rounds in which v is not part of a pair. For some arbitrary constant c, in the next $c \log \log n$ rounds v can increase its activity by at most a factor of $\sigma_{\oplus}^{c \log \log n} = \Theta(\operatorname{polylog} n)$. Let D_v be the random variable that counts the number of times v receives a message of one of its neighboring good leaders or heralds in those $c \log \log n$

rounds. In expectation, D_v increases at least every $\bar{m}$ rounds by one. With Chernoff we get

$$\mathbf{P}\left(D_v \leq \left(1 - \frac{1}{2}\right)\frac{c\log\log n}{\bar{m}}\right) \leq e^{-\frac{c\log\log n}{4\bar{m}}} \tag{27}$$

Recall that $|N_{r_i}| \leq \log^6 n$. We choose $c := 100\bar{m} \geq (20+6)\cdot 4\cdot \ln 2\cdot \bar{m}$. If $D_v > \frac{c\log\log n}{2\bar{m}}$ as in (27), then $\gamma(v)$ decreases enough in those $c\log\log n$ rounds to "touch" γ_{min} at least once. Now we can make a union bound over all nodes in N_{r_i} and we get that with probability $1 - \frac{1}{\log^{20} n}$ *all* nodes in N_{r_i} touch γ_{min} at least once in those $c\log\log n$ rounds.

Fourth part. Let us now count the amount of rounds r in which $I(u,r)$ does not hold. We keep the definition of rounds $r_1 < r_2 < \cdots < r_k$ and of events $\mathcal{E}_{r_i}$. Due to the definition of good pairs, there can be no more good heralds in $N^{\delta+1}(u)$ than there are good leaders in $N^{\delta+2}$. Once being part of a good pair, a node can only stop being good by being knocked out by an MIS node, which then prevents the creation of new leaders, heralds and MIS in its neighborhood. MIS nodes in distance $\delta+3$ can still influence good leaders and heralds in $N^{\delta+1}(u)$. However, no more than α good leaders (and their corresponding good heralds) and no more than α^2 good heralds (and their corresponding good leaders) can be knocked back to state $\mathbb{A}$ by an MIS node. Therefore, no more than $2\alpha^{\delta+5} = O(\sqrt{\log n})$ such events $\mathcal{E}_r$ can happen, i.e., $k = O(\sqrt{\log n})$.

We split the interval $[\bar{r}, \bar{r} + \tau_{progress}]$ into ℓ smaller intervals $(J_i)_{1\leq i\leq \ell}$ of length $c\log\log n$ each, i.e., $\ell = \frac{\tau_{progress}}{c\log\log n}$. Then we color each such interval J_i *red* if it contains one of the events $\mathcal{E}_r$ and we color J_i *orange* if its preceding interval J_{i-1} contains such an event, but not J_i itself. All other intervals are colored *blue*. From the third part we know the following. Independently of the activity values of all nodes in $N_r = N(W_r) \setminus W_r$ in some round r, if no event $\mathcal{E}_{r'}$ happens in round $r' < r + c\log\log n$, then with probability $1 - \log^{-20} n$ all those nodes touch γ_{min} within the next $c\log\log n$ rounds after round r. I.e., in every orange and blue interval, this is likely to happen at least once. If this happens in some orange or blue interval J_i, let us call this event $\mathcal{M}_i$. Any node that touches γ_{min} in interval J_i cannot recover its activity by a factor higher than

$$\sigma_{\oplus}^{c\log\log n} = 2^{\frac{24c\log\log n}{2000\bar{m}}} = \log^{\frac{24\cdot 100\bar{m}}{2000\bar{m}}} n = \log^{1.2} n,$$

until the end of interval J_{i+1}. In other words, if J_i is blue, then J_{i-1} cannot be red and if $\mathcal{M}_{i-1}$ happened, then $I(u,r)$ holds throughout the whole interval J_i. Let $\mathcal{I}$ be the index set for all blue intervals. Since $k = O(\sqrt{\log n})$, the total number of rounds in red and orange intervals are both in $o(\log n)$. Thus the number $\ell' := |\mathcal{I}|$ of blue intervals is in $(1-o(1))\ell$, i.e., $\ell' \geq (1-\varepsilon/2)\ell$. We define random variables X_i that evaluate to 1 if $\mathcal{M}_{i-1}$ does *not* hold and to 0 otherwise. Let $X := \sum_{i\in\mathcal{I}} X_i$, $p := \log^{-20} n$, $\varepsilon' := \varepsilon/2$ and $\tau_{progress} \geq \frac{c}{\varepsilon'(1-\varepsilon')}\log n$, then

$$\begin{aligned}
\mathbf{P}(X \geq \varepsilon'\ell') \;&\leq\; \binom{\ell'}{\varepsilon'\ell'} p^{\varepsilon'\ell'} \overset{\binom{n}{k}\leq\left(\frac{en}{k}\right)^k}{\leq} \left(\frac{e\ell'}{\varepsilon'\ell'}\right)^{\varepsilon'\ell'} \log^{-20\varepsilon'\ell'} n \\
&\leq\; 2^{((\log\frac{e}{\varepsilon'}) - 20\log\log n)\varepsilon'(1-\varepsilon')\frac{\tau_{progress}}{c\log\log n}} \leq n^{-19}.
\end{aligned}$$

Hence the number of intervals of length $c\log\log n$, in which no violation of $I(u,r)$ occurs, is w.h.p. at least $(1-\varepsilon)\ell$, which concludes the proof. □

Lemma 4.17. *Assume that (P) holds. Let t_u be a round in which a node $u \notin N^1(\mathbb{M})$ has a neighbor $u' \notin N^1(\mathbb{M})$ in the herald filter. Then, w.h.p., within $\tau_{progress} = O(\log^2 n/\mathcal{F} + \log n)$ rounds a node in $N^1(\{u,u'\})$ joins the MIS.*

Proof. The constant $\tau_{progress}$ is chosen by the end of this proof, and it is in $O(\log^2 n/\mathcal{F} + \log n)$. Initially we only set it to be large enough to comply with Lemma 4.14 and Lemma 4.16. Assume that the statement is not true, i.e., no node in $N^1(\{u,u'\})$ joins the MIS in the given time bounds.

Both nodes u and u' stay competitive, i.e., they stay in the states $\mathbb{A}$, $\mathbb{H}'$, $\mathbb{L}'$, $\mathbb{H}$ or $\mathbb{L}$, as states $\mathbb{M}$ and $\mathbb{E}$ would clearly imply the creation of an MIS node in $N^1(\{u,u'\})$. We now apply Lemma 4.14 to both nodes to get

$$\left|\left\{t' \in [t_u+1, t_u+\tau_{progress}] : \gamma(u,t') = \gamma(u',t') = 1/2\right\}\right| \quad \geq \quad (1-\varepsilon)\tau_{progress}. \tag{28}$$

We also apply Lemma 4.16 to extend (28) to

$$\left|\left\{t' \in [t_u+1, t_u+\tau_{progress}] : \gamma(u,t') = \gamma(u',t') = 1/2 \wedge I(u,t') \text{ holds}\right\}\right| \quad \geq \quad (1-2\varepsilon)\tau_{progress}. \tag{29}$$

We call rounds r in which $\gamma(u) = \gamma(u') = 1/2$ and $I(u,r)$ holds, a *promising round*. In such a round clearly $\Gamma(u) \geq 1$.

We first look at the case in which a round is promising, $\Gamma(u) < 5\alpha$ and u is $\frac{1}{5\alpha}$-fat. By Lemma 4.7 with probability $\Omega(\pi_\ell)$ within 7 rounds either a good pair or an MIS node is created nearby, an event that would contradict our assumption. Hence, as long as $\Theta(\pi_\ell^{-1}\log n)$ such rounds appear, w.h.p. such a contradicting event occurs. We choose $\tau_{progress}$ sufficiently large such that $\varepsilon\tau_{progress}$ such rounds would cause a contradiction.

W.l.o.g. we thus assume that within $\tau_{progress}$ rounds, a $(1-3\varepsilon)$ fraction of these are promising, but it does not hold that both $\Gamma(u) < 5\alpha$ and u is $\frac{1}{5\alpha}$-fat. For simplicity we exclude these cases from our definition of a promising round.

In all promising rounds, in distance δ there must be a node w, for which $\Gamma(w) \geq 5\alpha$ and w is η-fat. This is obvious if $\Gamma(u) \geq 5\alpha$, as then either u is already that node, or there exists a chain of activity sums, which increase in each step by a factor of at least η^{-1}. By definition of η, such a chain has length at most δ. If $\Gamma(u) < 5\alpha$, then u cannot be $\frac{1}{5\alpha}$-fat by our renewed definition of a promising round. But since $\Gamma(u) \geq 1$, u then must have a neighbor v for which $\Gamma(v) \geq 5\alpha$ and we can apply the same logic as before to find that η-fat node w.

For such a node w the requirements for Lemma 4.7 hold.

Now let $t_{1,prom}$ be the first promising round, $t_{2,prom}$ the second and so on.

Let $T_{1,prom}$ be the random variable that counts the number of promising rounds $t_{i,prom}$ until the first time a good pair or an MIS node is created in $N^\delta(u)$ or *has been* created in a non-promising round between $t_{i-1,prom}$ and $t_{i,prom}$ (we denote such an event as A). I.e.,

$$T_{1,prom} := \min\left\{i : \text{ A happens in } (t_{i-1,prom}, t_{i,prom}]\right\}.$$

$T_{j,prom}$ for $j > 1$ is defined accordingly. By Lemma 4.7 event A happens in expectation at least every $(c_1\pi_\ell)^{-1}$ promising rounds for some constant c_1 depending only on η.

Let M_u be the random variable that counts the number of events A happening in $N^\delta(u)$ until one happens in $N^1(\{u,u'\})$, i.e., until either a good herald or good leader is in $N^1(\{u,u'\})$. The number of good leaders and the number of MIS nodes that can co-exist in $N^\delta(u)$ is at most $2\alpha^\delta \leq 2\sqrt{\log n}$. If a node stops being a good leader, then it must have been knocked out by an MIS node, which still limits the total number of these events to $3\sqrt{\log n}$.

Note also that usually after an event A there is an $\Omega(\log\log n)$ pause until the next promising round happens, because it takes at least that many rounds for good pairs to reduce a neighbor's activity beneath the threshold γ_{low}, if that neighbors activity level is in $\Omega(1)$ by the time A happens. We account for that by "paying" $c_3 \log\log n$ rounds for each such event, for some constant c_3.

Now, using a Chernoff bound we get that $\sum_{i=1}^{M_u} T_{i,prom} \leq 6\sqrt{\log n}\left(\frac{1}{c_1\pi_\ell} + c_3\log\log n\right) + c_2\log n$, w.h.p., for some constant c_2. With suitably large chosen $\tau_{progress}$, this is less than $(1-3\varepsilon)\tau_{progress}$, causing the last contradiction to our assumption and thus finishing the proof. $\square$

Lemma 4.18. *W.h.p., property (P) is not violated throughout the runtime of the algorithm.*

Proof. Assume in round $\bar{t}$ it happens for the first time that 2 neighboring nodes u_1 and u_2 are in $\mathbb{M}$. W.l.o.g., we let u_1 enter the MIS before u_2 or at the same time and denote those times correspondingly $t_1 \leq t_2 = \bar{t}$. Additionally, for all pairs of nodes that violate (P) in round t_2 we choose the pair for which t_1, the 'age' of the older MIS node, is minimized, and ties among candidates for those positions are broken by *IDs*, first for u_1, then for u_2.

There are two ways for a node in $\mathbb{A}$ to join the MIS: either by increasing its *lonely* counter to $\tau_{\textit{lonely}}$ or by passing through stages $\mathbb{L}'$ and $\mathbb{L}$.

(A) Let us first assume that both nodes u_1 and u_2 enter state $\mathbb{M}$ by finishing their respective red-blue games. If $t_1 < t_2$, then u_1 started its first red-blue game before u_2 did, but by Lemma 4.10, they cannot be in state $\mathbb{L}$ at the same time. Thus, u_1 joins the MIS before u_2 becomes a leader. But as a leader, u_1 disrupts handshake channel $\mathcal{H}$ every second round, and then as an MIS node, it does so at least every second round, preventing u_2 from ever becoming a leader. Therefore, let $t_1 = t_2$. Then u_1 and u_2 were neighboring leaders for $\tau_{\textit{red-blue}}/8 = \Omega(\log n)$ red-blue games, and in each such game having a probability of at least $1/2$ to conflict with each others' red-blue game. By Chernoff, w.h.p., this is not possible.

(B) Let us next assume that both nodes enter state $\mathbb{M}$ through their *lonely* counters. First assume that $t_2 - t_1 \geq \tau_{\textit{lonely}}/2 \geq \tau_{\textit{notification}}$. But then in $[t_1, t_2 - 1]$ (P) holds true and with Lemma 4.13, u_2 gets eliminated by u_1 before it can become an MIS node. Thus, let $t_2 - t_1 < \tau_{\textit{lonely}}/2$. But then during rounds $[t_1 - \tau_{\textit{lonely}}/2, t_1]$ both nodes are in the active states $\mathbb{A}$, $\mathbb{H}'$, $\mathbb{L}'$, $\mathbb{H}$ and $\mathbb{L}$ and they do not neighbor an MIS node. The latter stems from the following. If u_1 would neighbor an MIS node u_3, then the first violation of (P) would happen in round t_1, a contradiction if $t_1 < t_2$. For $t_1 = t_2$, nodes u_1 and u_2 contradict our choice of the pair being investigated, as u_3 would have had a higher age. Similar argumentation keeps u_2 from neighboring an MIS node. Hence, assume that none of both nodes neighbor an MIS node. Lemma 4.17 tells us that by round $t_1 - \tau_{\textit{lonely}} + \tau_{\textit{progress}} \ll t_1$, an MIS node v would arise in $N^1(\{u_1, u_2\})$, with $v \neq u_1, u_2$. The remaining rounds, by Lemma 4.13 suffice for v to eliminate u_1 or u_2 w.h.p., again contradicting our initial assumption.

(C) Now let u_1 join the MIS via loneliness and u_2 by being a leader. If $t_2 - t_1 \geq \tau_{\textit{red-blue}}/2 = \Omega(\log n)$, then in each 8-round red-blue game after round $t_1 + 1$, u_1 has a constant chance to disrupt u_2's game on channel $\mathcal{G}$. Choosing $\tau_{\textit{red-blue}}$ large enough guarantees us that, w.h.p., this cannot happen. If $t_2 - t_1 < \tau_{\textit{red-blue}}/2$, then in each 8-round red-blue game in $(t_2 - \tau_{\textit{red-blue}}, t_1)$, there is a constant probability that u_2 transmits on a report channel $\mathcal{R}_i$ on which no other neighbor of u_1 transmits. The argumentation is similar to the one in Lemma 4.16, except that u_1 does not neighbor an MIS node yet (see argumentation above w.r.t. to our choice of u_1 and u_2): There are at most $2\alpha^2$ good leaders and heralds neighboring u_1, and *all* bad leaders/heralds get knocked out with constant probability in each red-blue game, see Lemma 4.12. At the same time u_1 listens on the same report channel $\mathcal{R}_i$ with probability at least $\frac{1}{2n_{\mathcal{R}}} = \Omega(1)$. Therefore, with $\tau_{\textit{red-blue}}$ large enough, w.h.p. u_1 hears from u_2 before round t_1 and thus resets its *lonely* counter.

(D) Last switch u_1's and u_2's roles from the previous case. In this case, u_2 is neighboring u_1 throughout its leadership state, and with analogous argumentation, regardless whether $u_1 \in \mathbb{L}$ or $u_1 \in \mathbb{M}$, u_2 hears from u_1 with high probability. $\square$

Now we have everything at hand to prove Theorem 3.1.

Proof of Theorem 3.1. As stated earlier, Lemma 4.2 provides that the runtime of the decay filter is within $O\big(\log^2 n/\mathcal{F} + \log n\big)$, i.e., for a node u executing Algorithm 2 (the decay filter) by that time one node $v \in N_G^1(u)$ enters the herald filter. From Lemma 4.1 we get that over the course of $O(\log^2 n)$ rounds the maximum degree of the graph G' induced by all nodes in the herald filter is at most $O(\log^3 n)$.

Let thus u be a node that enters the herald filter. If it stays lonely for $\tau_{\mathit{lonely}} = \Theta\big(\log^2 n/\mathcal{F} + \log n\big)$ rounds, then u joins the MIS and we are done. Hence assume that u does hear from a neighboring node u' in the herald filter before τ_{lonely} rounds have passed. We can now apply Lemma 4.17 to get an MIS node v created within $\tau_{\mathit{progress}} = O\big(\log^2 n/\mathcal{F} + \log n\big)$ rounds. It either neighbors u, in which case within $\tau_{\mathit{notification}} = O(\log n)$ rounds u is decided w.h.p., or it neighbors u', which is also then eliminated in $\tau_{\mathit{notification}}$ rounds. That way u can become lonely again. However, since an MIS node has been created in $N^2(u)$, this can happen at most α^2 times. Thus, at most $\tau_{\mathit{runtime}} = 2\alpha^2\tau_{\mathit{lonely}}$ rounds after u entered the herald filter, u is decided. □

References

[1] I. 802.11. Wireless LAN MAC and Physical Layer Specifications, March 2012.

[2] Z. Alliance. Zigbee specification. *ZigBee Document 053474r06*, 1, 2005.

[3] N. Alon, L. Babai, and A. Itai. A Fast and Simple Randomized Parallel Algorithm for the Maximal Independent Set Problem. *Journal of Algorithms*, 1986.

[4] R. Bar-Yehuda, O. Goldreich, and A. Itai. On the Time-Complexity of Broadcast in Multi-Hop Radio Networks: An Exponential Gap Between Determinism and Randomization. *Journal of Computer and System Sciences*, 45(1):104–126, 1992.

[5] Bluetooth Consortium. *Bluetooth Specification Version 4.2*, December 2014.

[6] K. Censor-Hillel, S. Gilbert, F. Kuhn, N. Lynch, and C. Newport. Structuring Unreliable Radio Networks. In *Proc. ACM Symp. on Principles of Distributed Computing (PODC)*, pages 79–88, 2011.

[7] I. Chlamtac and S. Kutten. On Broadcasting in Radio Networks–Problem Analysis and Protocol Design. *IEEE Transactions on Communications*, 33(12):1240–1246, 1985.

[8] S. Daum, M. Ghaffari, S. Gilbert, F. Kuhn, and C. Newport. Maximal Independent Sets in Multichannel Radio Networks. In *Proc. ACM Symp. on Principles of Distributed Computing (PODC)*, 2013.

[9] S. Daum, M. Ghaffari, S. Gilbert, F. Kuhn, and C. Newport. Maximal Independent Sets in Multichannel Radio Networks. Technical Report 275, University of Freiburg, Dept. of Computer Science, 2013.

[10] S. Daum, S. Gilbert, F. Kuhn, and C. Newport. Leader Election in Shared Spectrum Networks. In *Proc. ACM Symp. on Principles of Distributed Computing (PODC)*, 2012.

[11] S. Daum, F. Kuhn, and C. Newport. Efficient Symmetry Breaking in Multi-Channel Radio Networks. In *Proc. Int. Symp. on Distributed Computing (DISC)*, 2012.

[12] S. Daum, F. Kuhn, and C. Newport. Efficient Symmetry Breaking in Multi-Channel Radio Networks. Technical Report 271, University of Freiburg, Dept. of Computer Science, 2012.

[13] S. Dolev, S. Gilbert, R. Guerraoui, D. R. Kowalski, C. Newport, F. Kuhn, and N. Lynch. Reliable Distributed Computing on Unreliable Radio Channels. In *Proceedings of the 2009 MobiHoc S^3 Workshop*, 2009.

[14] S. Dolev, S. Gilbert, R. Guerraoui, F. Kuhn, and C. Newport. The wireless synchronization problem. In *Proc. ACM Symp. on Principles of Distributed Computing (PODC)*, pages 190–199, 2009.

[15] S. Dolev, S. Gilbert, R. Guerraoui, and C. Newport. Gossiping in a Multi-Channel Radio Network: An Oblivious Approach to Coping with Malicious Interference. In *Proc. Int. Symp. on Distributed Computing (DISC)*, 2007.

[16] S. Dolev, S. Gilbert, R. Guerraoui, and C. Newport. Secure Communication Over Radio Channels. In *Proc. ACM Symp. on Principles of Distributed Computing (PODC)*, 2008.

[17] S. Dolev, S. Gilbert, M. Khabbazian, and C. Newport. Leveraging Channel Diversity to Gain Efficiency and Robustness for Wireless Broadcast. In *Proc. Int. Symp. on Distributed Computing (DISC)*, 2011.

[18] A. Ephremides, J. E. Wieselthier, and D. J. Baker. A Design Concept for Reliable Mobile Radio Networks with Frequency Hopping Signaling. *Proceedings of the IEEE*, 75(56–73), 1987.

[19] S. Gilbert, R. Guerraoui, D. Kowalski, and C. Newport. Interference-Resilient Information Exchange. In *Proc. IEEE Int. Conf. on Computer Communications (INFOCOM)*, 2009.

[20] F. Kuhn, T. Moscibroda, T. Nieberg, and R. Wattenhofer. Fast Deterministic Distributed Maximal Independent Set Computation on Growth-Bounded Graphs. In *Proc. Int. Symp. on Distributed Computing (DISC)*, 2005.

[21] M. Luby. A Simple Parallel Algorithm for the Maximal Independent Set Problem. *SIAM Journal on Computing*, 15(4):1036–1053, 1986.

[22] T. Moscibroda and R. Wattenhofer. Maximal independent sets in radio networks. In *Proc. ACM Symp. on Principles of Distributed Computing (PODC)*, pages 148–157, 2005.

[23] C. Newport. Radio Network Lower Bounds Made Easy. In *Proceedings of the International Symposium on Distributed Computing (DISC)*, 2014.

[24] S. Schmid and R. Wattenhofer. Algorithmic models for sensor networks. In *Proc. Int. Workshop on Parallel and Distributed Real-Time Systmes (WPDRTS)*, pages 1–11, 2006.

[25] J. Schneider and R. Wattenhofer. A Log-Star Distributed Maximal Independent Set Algorithm for Growth-Bounded Graphs. In *Proc. ACM Symp. on Principles of Distributed Computing (PODC)*, pages 35–44, 2008.

[26] M. Sherman, A. Mody, R. Martinez, C. Rodriguez, and R. Reddy. IEEE Standards Supporting Cognitive Radio and Networks, Dynamic Spectrum Access, and Coexistence. *IEEE Communications Magazine*, 46(7):72–79, 2008.

[27] M. Strasser, C. Pöpper, and S. Capkun. Efficient Uncoordinated FHSS Anti-jamming Communication. In *Proc. ACM Int. Symp. on Mobile Ad Hoc Networking and Computing (MOBIHOC)*, 2009.

[28] M. Strasser, C. Pöpper, S. Capkun, and M. Cagalj. Jamming-resistant Key Establishment using Uncoordinated Frequency Hopping. In *Proc. of the IEEE Symposium on Security and Privacy (S & P)*, 2008.

[29] P.-J. Wan, K. M. Alzoubi, and O. Frieder. Distributed Construction of Connected Dominating Set in Wireless Ad Hoc Networks. In *Proceedings of the IEEE International Conference on Computer Communications (INFOCOM)*, 2002.

A Properties of Graphs with Bounded Independence

We need a few statements about bounded independence graphs. The proofs are provided in [9].

Lemma A.1. *Let $G = (V, E)$ be a graph and assume that every node $u \in V$ has a positive weight $w_u > 0$. Define $W := \sum_{v \in V} w_v$ and for each $u \in V$, $W_u := \sum_{v \in N_G^1(u)} w_v$. It then holds that*

$$\sum_{v \in V} \frac{w_v}{W_v} \leq \alpha(G) \quad \textit{and} \tag{30}$$

$$\sum_{v \in V} w_v \cdot W_v \geq \frac{W^2}{\alpha(G)}, \tag{31}$$

where $\alpha(G)$ is the independence number of G.

Lemma A.2. *Let $G = (V, E)$ be a graph and assume that every node $u \in V$ has a positive weight $w_u > 0$. Define $W := \sum_{v \in V} w_v$ and for each $u \in V$, $W_u := \sum_{v \in N_G^+(u)} w_v$. Let $V_{heavy} \subseteq V$ be the set of nodes v for which $W_v \geq \frac{W}{2\alpha(G)}$. The total weight of nodes in V_{heavy} is at least*

$$\sum_{v \in V_{heavy}} w_v > \frac{W}{2\alpha(G)}.$$

B Candidate Election—Statements and Proofs for $\mathbb{A}$ (and $\mathbb{L}'$)

B.1 Lemma 4.6, "W.h.p. nothing happens"

Let k be a positive integer *constant*, r some round, u some node in the herald filter in round r, i an index from $1, \ldots, n_{\mathcal{A}}$, S be a (possibly empty) subset of $N^k(u) \cap \mathbb{A}$. Furthermore let $\partial S \subset S$ be *the* subset of S that has connections outside of S, but in $N^k(u)$, i.e., $\partial S := S \cap N\big(N^k(u) \setminus S\big)$. At last, let $S = S^n \dot\cup S^b \dot\cup S^l$ be a partition of S. We call the tuple $(k, u, r, i, S))$ a *constellation*. For a constellation the following events are defined for round r:

- $\mathcal{S}_{\neg i}/\mathcal{S}^n_{\neg i}$: *no* node in S/S^n operates on $\mathcal{A}_i$,
- $\mathcal{S}_i/\mathcal{S}^b_i/\mathcal{S}^l_i$: *all* nodes in $S/S^b/S^l$ operate on $\mathcal{A}_i$,
- $\partial\mathcal{S}_{\neg i}$: *no* node in ∂S operates on $\mathcal{A}_i$,
- H_i: *no* node in $N^k(u) \setminus S$ receives a message on channel $\mathcal{A}_i$,
- $H_{\neg i}$: *no* node in $N^k(u)$ receives a message on some channel $\mathcal{A}_j \neq \mathcal{A}_i$ and
- H: *no* node in $N^k(u)$ receives a message on any channel $\mathcal{A}_1, \ldots, \mathcal{A}_{n_{\mathcal{A}}}$.

Lemma B.1. *Let (k, u, r, i, S) be a constellation. Then,*

(1) $\mathbf{P}(H) = 1 - O(\pi_\ell \alpha^k)$

(2) $\mathbf{P}(H_{\neg i} | \mathcal{S}_{\neg i}) = 1 - O(\pi_\ell \alpha^k)$

(3) $\mathbf{P}(H_{\neg i} | \mathcal{S}_i) = 1 - O(\pi_\ell \alpha^k)$

(4) $\mathbf{P}(H_{\neg i} | \mathcal{S}^n_{\neg i} \wedge \mathcal{S}^b_i \wedge \mathcal{S}^l_i) = 1 - O(\pi_\ell \alpha^k)$

(5) $\mathbf{P}(H_i | \partial\mathcal{S}^n_{\neg i}) = 1 - O(\pi_\ell \alpha^k)$

Look at (1). The lemma says that the probability for a herald candidate to be created in any single round for any neighborhood of constant radius is at most linear in π_ℓ. Since π_ℓ is an arbitrarily small constant parameter chosen by us, we can make the probability for this event arbitrarily small. The proof for (1)

is exactly the same as the proof for Lemma 8.6 in [9], with $\alpha(k)$ replaced by α^k. We provide its proof nevertheless, as (2) and (3) are new results that do directly depend on (1). What (2) says, is, that even if we condition on some nodes $S \subset N^k(u)$ *not* to operate on channel $\mathcal{A}_i$, this does not increase the chance (significantly—i.e., by more than a constant factor) for these or other nodes to receive anything on some channel $\mathcal{A}_j \neq \mathcal{A}_i$. Analogously, (3) claims that if we condition on some nodes $S \subset N^k(u)$ to *definitely* operate on channel $\mathcal{A}_i$, this does still not affect the probability for any other node to receive any message on some other channel $\mathcal{A}_j \neq \mathcal{A}_i$. (4) is a combination of (2) and (3), plus we even fix the knowledge of which nodes, that operate on $\mathcal{A}_i$, do broadcast (S^b) or listen (S^l). (5) has already been proven in [9] as Claim 8.9, except that there k was fixed to 2. The analogous proof is provided as well for sake of completeness.

Proof. We first prove (1). For the whole proof we only use the graph $G_{\mathbb{A}}$ induced by nodes in state $\mathbb{A}$ in round r. We also solely focus on nodes v that do have at least one active neighbor in $G_{\mathbb{A}}$, as isolated nodes cannot become herald candidates. We will use the notation $N_{\mathbb{A}}(v)$ and $N_{\mathbb{A}}^d(v)$ to refer to $N_{G_{\mathbb{A}}}(v)$ and $N_{G_{\mathbb{A}}}^d(v)$, respectively.

To become a herald candidate, a node v in state $\mathbb{A}$ must receive a message from one of its neighbors on one of the channels $\mathcal{A}_1, \ldots, \mathcal{A}_{n_{\mathcal{A}}}$. This is only possible if in round r, v chooses to listen on a channel $\mathcal{A}_j$ and exactly one of v's neighbors in $G_{\mathbb{A}}$ broadcasts on channel $\mathcal{A}_j$.

Consider an arbitrary channel $\mathcal{A}_j$ from the herald election channels $\mathcal{A}_1, \ldots, \mathcal{A}_{n_{\mathcal{A}}}$. Let $p_v(j) = 2^{-j} \cdot \gamma(v)$ be the probability that an active node v chooses to broadcast or listen on channel $\mathcal{A}_j$. In addition, we define $P_v(j) := 2^{-j}\Gamma(v) = \sum_{w \in N_{\mathbb{A}}^1(v)} p_w(j)$. Let $B_j^{v,w}$ be the event that v listens on channel $\mathcal{A}_j$, while exactly one of its neighbors $w \in N_{\mathbb{A}}(v)$ transmits on channel $\mathcal{A}_j$ and all other neighbors $w' \in N_{\mathbb{A}}(v)$ are either not on channel $\mathcal{A}_j$ or they choose to listen as well.

$$
\begin{aligned}
\mathbf{P}(B_j^{v,w}) &= \pi_\ell p_v(j) \cdot (1-\pi_\ell) p_w(j) \cdot \prod_{w' \in N_{\mathbb{A}}(v) \setminus \{w\}} \big(1 - p_{w'}(j)(1-\pi_\ell)\big) \\
&\le \pi_\ell p_v(j) p_w(j) \cdot \prod_{w' \in \{v,w\}} \frac{1}{1-(1-\pi_\ell)p_{w'}(j)} \cdot \prod_{w' \in N_{\mathbb{A}}^1(v)} \big(1 - p_{w'}(j)(1-\pi_\ell)\big) \\
&\le \pi_\ell p_v(j) p_w(j) \cdot 4 \cdot e^{-\frac{1}{2}P_v(j)} \\
&= \pi_\ell 2^{-2j} \gamma(v)\gamma(w) \cdot 4 e^{-2^{1-j}\Gamma(v)}
\end{aligned}
\tag{32}
$$

In the last inequality, we use that $p_{w'}(j) \le \frac{1}{2}$ and that $\pi_\ell \le \frac{1}{2}$. Define B_j^v to be the event that v listens on $\mathcal{A}_j$ and exactly one of its neighbors transmits on that channel. Since $B_j^v = \bigcup_{w \in N_{\mathbb{A}}(v)} B_j^{v,w}$ and the events $B_j^{v,w}$ are disjoint for different w, we have

$$
\mathbf{P}(B_j^v) = \sum_{w \in N_{\mathbb{A}}(v)} \mathbf{P}(B_j^{v,w}) \le \pi_\ell p_v(j) P_v(j) \cdot 4 e^{-\frac{1}{2}P_v(j)} =: C_j^v.
$$

For any $x > 0$ and constant c, $cx^2 e^{-x} = O(1)$, which by using $x = P_v(j)$ implies that $C_j^v = O\big(\pi_\ell \frac{p_v(j)}{P_v(j)}\big) = O\big(\pi_\ell \frac{\gamma(v)}{\Gamma(v)}\big)$ for any fixed j. Next we show that $\sum_{j=1}^{n_{\mathcal{A}}} C_j^v = O\big(\pi_\ell \frac{\gamma(v)}{\Gamma(v)}\big)$, too.

$$
\frac{C_{j+1}^v}{C_j^v} = \frac{p_v(j+1)}{p_v(j)} \frac{P_v(j+1)}{P_v(j)} e^{-\frac{1}{2}P_v(j+1)+\frac{1}{2}P_v(j)} = \frac{1}{4} e^{\frac{1}{4}\Gamma(v)2^{-j}} < \frac{1}{2} \qquad \forall j \ge \log \Gamma(v) \tag{33}
$$

$$
\frac{C_j^v}{C_{j+1}^v} = \frac{p_v(j)}{p_v(j+1)} \frac{P_v(j)}{P_v(j+1)} e^{-\frac{1}{2}P_v(j)+\frac{1}{2}P_v(j+1)} = 4 e^{-\frac{1}{4}\Gamma(v)2^{-j}} < \frac{1}{2} \qquad \forall j \le \log \Gamma(v) - 4 \tag{34}
$$

We can therefore deduce the upper bounds

$$\sum_{j\geq\log\Gamma(v)} C_j^v \leq 2C_{\lceil\log\Gamma(v)\rceil}^v \text{ and } \sum_{j\leq\log\Gamma(v)} C_j^v \leq 2C_{\max\{1,\lfloor\log\Gamma(v)-4\rfloor\}}^v,$$

proving the claim that $\sum_{j=1}^{n_\mathcal{A}} C_j^v = O\big(\pi_\ell \frac{\gamma(v)}{\Gamma(v)}\big)$.

Using Lemma A.1, choosing $G' := G_\mathbb{A}[N_\mathbb{A}^k(u)]$, $w(v) := \gamma(v)$ and $W(v) := \Gamma(v)$, we get that $\sum_{v\in G'} \frac{\gamma(v)}{\Gamma(v)} \leq \alpha(G') \leq \alpha^k$. (Note that the independence number of a graph is larger than or equal to the independence number of any induced subgraph.)

Let B^v be the event that v moves from state $\mathbb{A}$ to $\mathbb{H}'$ and $B = \bigcup_{v\in N^k(u)} = \bigcup_{v\in N_\mathbb{A}^k(u)}$. Then,

$$\mathbf{P}(B) \leq \sum_{v\in N_\mathbb{A}^k(u)} \mathbf{P}(B^v) \leq \sum_{v\in N_\mathbb{A}^k(u)} \sum_{j=1}^{n_\mathcal{A}} C_j^v = \sum_{v\in N_\mathbb{A}^k(u)} O\left(\pi_\ell \frac{\gamma(v)}{\Gamma(v)}\right) = O\left(\pi_\ell \alpha^k\right).$$

Choosing a sufficiently small π_ℓ concludes the proof of (1).

For (2) let a proper set S be given. We apply the same calculations as before but we condition now on S_i, i.e., that all nodes in S do *not* operate on channel $\mathcal{A}_i$. We let $p_v(j)$ still be the probability that node v operates on channel $\mathcal{A}_j$, but its value might now be different. If $v \notin S$, then this probability is exactly $p_v(j)$ as above, i.e., equals $2^{-j}\gamma(v)$. Otherwise, this probability is larger than that by a factor of $\frac{1}{1-2^{-i}}$ due to conditioning on v not operating on $\mathcal{A}_i$. This term maximizes at $i = 1$ and then evaluates to 2, i.e., $p_v(j) \leq 2\cdot 2^{-j}\gamma(v)$. The same thought needs to be applied to all nodes and therefore also $P_v(j)$ can change: It now ranges between $2^{-j}\Gamma(v)$ and $2^{1-j}\Gamma(v)$. Since we are looking for an upper bound in Equation 32, we use that $e^{-\frac{1}{2}P_v(j)} \leq e^{-2^{1-j}\Gamma(v)}$. In total we get that the right hand side of Equation 32 increases by a factor at most 4.

The remainder of the proof is completely analogous to the one for (1), except that in any summations over all channels we omit channel $\mathcal{A}_i$ (resp. index i), which just benefits the cause.

For (3) we assume again that a proper set S is given and we condition on $\tilde{S}_i$, i.e., all nodes in S *do* operate on channel $\mathcal{A}_i$. But we are only interested in heralds being created on channels $\mathcal{A}_j \neq \mathcal{A}_i$. We apply the following simple adaption to our initial setup. We only focus on the graph $G'_\mathbb{A}$ induced by nodes in state $\mathbb{A}$ that will *not* operate on channel $\mathcal{A}_i$ in round r, i.e., we exclude all nodes from S, but also nodes outside of S whose local random bits indicate that they choose to operate on $\mathcal{A}_i$ this round. In this case all calculations stay *exactly* the same, except that channel $\mathcal{A}_i$ (resp. index i) needs to be removed from all summations, unions (as a matter of fact, the probability of any analyzed event referring to channel $\mathcal{A}_i$ equals zero in $G'_\mathbb{A}$) and observations in general. When analyzing values C_j^v the index i needs to be skipped. Finally, the subgraph G' used in the last step must be based on $G'_\mathbb{A}$.

For (4) let initially $S^n = \emptyset$. Then the statement is the same as in (3), except that we not only know that nodes in S operate on channel $\mathcal{A}_i$, we also know whether they broadcast (set S^b) or listen (set S^l). However, this additional knowledge does not affect the proof of (3)—what happens among nodes on other channels than $\mathcal{A}_i$ is completely detached from events on $\mathcal{A}_i$. Let now $S^n \neq \emptyset$, then (4) is a simple combination of (2) and this expanded version of (3).

The proof of (5) is exactly the same as the proof for Claim 8.9 in [8], which is equivalent to setting $k = 2$.

In the following, we use the notation $\mathcal{N} := N_\mathbb{A}^k(u, r) \setminus S$. For a node $x \in \mathcal{N}$, let B_x be the event that node x receives a message on channel $\mathcal{A}_i$. Event B_x occurs iff x listens on channel $\mathcal{A}_i$ and exactly one of its neighbors broadcasts on channel $\mathcal{A}_i$. The probability for a node $x \in \mathcal{N}$ to pick channel $\mathcal{A}_i$ is $\gamma(x)\cdot 2^{-i}$.

We therefore have

$$\begin{aligned}\mathbf{P}(B_x|\partial\mathcal{S}_{\neg i}) &= \frac{\pi_\ell\gamma(x,r)}{2^i}\sum_{z\in N_{\mathbb{A}}(x)\setminus S}\frac{(1-\pi_\ell)\gamma(z,r)}{2^i}\cdot\prod_{y\in N_{\mathbb{A}}(x)\setminus(S\cup\{z\})}\left(1-\frac{(1-\pi_\ell)\gamma(y,r)}{2^i}\right)\\ &\leq \frac{\pi_\ell\gamma(x,r)}{2^i}\frac{\Gamma(x,r)}{2^i}\cdot e^{-\Gamma(x,r)2^{-i}} = O\left(\frac{\gamma(x,r)}{\Gamma(x,r)}\cdot\pi_\ell\right).\end{aligned}$$

Let X be the number of nodes $x\in\mathcal{N}$ that receive a message on channel $\mathcal{A}_i$ in round r. For the expectation of X, we then get

$$\mathbb{E}[X|\partial\mathcal{S}_{\neg i}] = O(\pi_\ell)\cdot\sum_{x\in\mathcal{N}}\frac{\gamma(x,r)}{\Gamma(x,r)} = O(\pi_\ell).$$

The second equation follows from Lemma A.2 because the graph induced by $\mathcal{N}$ has independence at most α^k. Applying the Markov inequality, we get $\mathbf{P}(X\geq 1|\partial\mathcal{S}_{\neg i})\leq\mathbb{E}[X|\partial\mathcal{S}_{\neg i}] = O(\pi_\ell)$, which concludes the proof of (5). □

B.2 Lemma 4.7

Lemma B.2. *Let t be a round in which for a node u in state $\mathbb{A}$ in the herald filter the following holds:*

- *there is no herald candidate in $N^2(u)$,*
- *all nodes $v\in N^2(u)$ that neighbor a herald or leader, have $\gamma(v)\leq\gamma_{low}$,*
- *all nodes in $N^2(u)$ neighboring MIS nodes are eliminated,*
- $\Gamma(u)\geq 1$,

If in addition it holds that either
(a) $\Gamma(u)<5\alpha$, *u is $\frac{1}{5\alpha}$-fat and $\gamma(u)=\frac{1}{2}$, or*
(b) $\Gamma(u)\geq 5\alpha$ *and u is η-fat.*
Then by the end of round $t'\in[t,t+7]$, with probability $\Omega(\pi_\ell)$ either a node in $N^2(u)$ joins the MIS or a good pair $(l,h)\in(\mathbb{L}\cap N^1(u))\times(\mathbb{H}\cap N^1(u))$ is created.

This proof is an adaption to the following lemma, from [9]:

Lemma B.3. *Let t be a round in which for a node u in state $\mathbb{A}$ in the herald filter it holds that there is no herald, leader, or herald candidate in $N^2(u)$. Furthermore, all neighbors of MIS nodes in $N^2(u)$ are in state $\mathbb{E}$, $\Gamma(u)\geq 1$, and either*
(a) $\Gamma(u)<3\alpha(1)$, *u is $\frac{1}{3\alpha(1)}$-fat, and $\gamma(u)=\frac{1}{2}$, or*
(b) *u is $\frac{1}{2}$-fat and $\Gamma(u)\geq 3\alpha(1)$.*
Then by round $t'\in[t,t+7]$, with probability $\Omega(\pi_\ell)$ either a node in $N^2(u)$ joins the MIS or a pair $(l,h)\in\mathbb{L}\times\mathbb{H}$ is created in $N^1(u)$ such that $(N(\{l,h\})\setminus\{l,h\})\cap(\mathbb{H}'\cup\mathbb{H}\cup\mathbb{L})=\emptyset$.

Note that we do not exclude the existence of leaders and heralds anymore. Instead we require such nodes to have 'silenced' their neighborhood, i.e., to have drove back such nodes' activity values. Good pairs are supposed to do so with a good probability, and while bad pairs can manage to do that, too (with a low probability), it does not affect the lemma's correctness if they do. Last, the probability for a new isolated leader-herald pair goes down a bit, but is still in $\Omega(\pi_\ell)$, yet the hidden constant now depends heavily on η. One of the main reason to change 3α from the original lemma to 5α was to ease argumentation in one part and to differentiate it from $n_{\mathcal{R}}\geq 3\alpha^2$—the value in this proof is completely unrelated to $n_{\mathcal{R}}$.

The proof follows mostly the lines of the one in [9], but in that paper a constant number of factors of $1/2$ accumulate, while here they have to be exchanged by factors of η—but this only causes changes in asymptotically irrelevant constants. In addition, nodes neighboring leaders or heralds need special attention. We give a full detailed proof here, despite its similarity to the proof in [9].

Proof. We make use of the notation $X(u,t)$ to indicate the value of local variable X at node u in round t. For the remainder of the proof, assume that in rounds $t,\dots,t+7$, no node in $N^2(u)$ joins the MIS, as otherwise, the claim of the lemma is trivially satisfied. In order to prove the lemma, we first show that either in round t or in round $t+1$, w.c.p., a herald candidate is created in $N^1(u)$. Formally, we define the event H_u as follows. In round t', event H_u occurs iff there are two neighboring nodes $v,w \in N^1_{\mathbb{A}}(u,t') \setminus N(\mathbb{L} \cup \mathbb{H})$ such that

- v and w both operate on a channel $\lambda \in \{\mathcal{A}_1,\dots,\mathcal{A}_{n_{\mathcal{A}}}\}$
- no other neighbor of v and w chooses channel λ, and
- no other node in $N^2_{\mathbb{A}}(u,t')$ receives a message on channel λ.

Clearly, if event H_u holds either in round t or $t+1$, the nodes v,w have a probability of $2\pi_\ell(1-\pi_\ell)$ of becoming a herald-leader candidate pair and no other herald candidate is created on channel λ in that round. Combined with appropriate applications of Lemma B.1.(1), this suffices to prove the claim of the lemma.

For a node $v \in N^1_{\mathbb{A}}(u,t')$, let $\Gamma^u_{\mathbb{A}}(v,t') := \sum_{w \in N^1(v,t') \cap N^1_{\mathbb{A}}(u,t')} \gamma(w,t')$ be the total activity value of all active nodes in round t' in the 1-neighborhood of v restricted to the 1-neighborhood of u. To estimate the probability that H_u occurs in a round $t' \in \{t,t+1\}$, we first show that in one of the two rounds $t' \in \{t,t+1\}$, with probability at least $\frac{1}{4}$ it holds that u is in state $\mathbb{A}$ and $\Gamma_{\mathbb{A}}(u,t') := \sum_{v \in N^1_{\mathbb{A}}(u,t')} \gamma(v,t') \geq \frac{3}{5} \cdot \Gamma(u,t)$; i.e., in one of both rounds, we have a high activity mass provided by nodes in state $\mathbb{A}$ as opposed to those in state $\mathbb{L}'$. Assume that the claim is not true for $t'=t$. As the lemma statement is based on the assumption that u is in state $\mathbb{A}$ in round t, this implies that $\Gamma_{\mathbb{A}}(u,t) < \frac{3}{5}\Gamma(u,t)$. Also by the assumptions of the lemma, in round t, no nodes in $N(u)$ are in states $\mathbb{H}'$, and those in $\mathbb{H}$ or $\mathbb{L}$ have an activity of less than γ_{low}. Since $|N^2(u) \cap N(\mathbb{L} \cup \mathbb{H})| \leq \Delta^2_{max} \leq \log^{2\kappa_\Delta} n \ll \gamma^{-1}_{low}$, we can assume that the total activity mass of nodes in $N^2(u) \cap N(\mathbb{L} \cup \mathbb{H})$ is less than $1/100$. As nodes w in states $\mathbb{M}$ and $\mathbb{E}$ have $\gamma(w)=0$ and thus do not contribute to $\Gamma(u)$, we therefore have

$$\Gamma_{\mathbb{L}'}(u,t) := \sum_{v \in N^1_{\mathbb{L}'}(u,t)} \gamma(v,t) \geq \Gamma(u,t) - \Gamma_{\mathbb{A}}(u,t) - \frac{1}{100} \overset{\Gamma(u,t)\geq 1}{\geq} \frac{99}{100}\Gamma(u,t) - \Gamma_{\mathbb{A}}(u,t).$$

Because by assumption, there are no nodes in state $\mathbb{H}'$ in round t, all nodes that are in state $\mathbb{L}'$ in round t switch back to state $\mathbb{A}$ for the next round. As by assumption, no nodes switch to states $\mathbb{M}$ or $\mathbb{E}$, and a node v that is in state $\mathbb{A}$ in round t can only move out of $\mathbb{A}$ if it decides to operate on one of the channels $\mathcal{A}_1,\dots,\mathcal{A}_{n_{\mathcal{A}}}$. This happens with probability at most $\gamma(v,t) \leq \frac{1}{2}$. Therefore, with probability at least $\frac{1}{2}$, at least half of the total activity value of the nodes in $N^1_{\mathbb{A}}(u,t)$ remains in state $\mathbb{A}$ for round $t+1$. And (independently) with probability at least $\frac{1}{2}$, also node u remains in state $\mathbb{A}$ for round $t+1$. Thus, with probability at least $\frac{1}{4}$, u is in state $\mathbb{A}$ in round $t+1$ and at least half of the total activity contributing to $\Gamma_{\mathbb{A}}(u,t)$ also contributes to $\Gamma_{\mathbb{A}}(u,t+1)$. Therefore, with probability at least $\frac{1}{4}$,

$$\begin{aligned}
\Gamma_{\mathbb{A}}(u,t+1) &\overset{(*)}{\geq} \left(\frac{98}{100}\Gamma(u,t) - \Gamma_{\mathbb{A}}(u,t)\right) + \frac{1}{2}\cdot\Gamma_{\mathbb{A}}(u,t) \\
&\overset{\Gamma(u,t)\geq 1}{\geq} \frac{98}{100}\Gamma(u,t) - \frac{1}{2}\cdot\Gamma_{\mathbb{A}}(u,t) \geq \frac{3}{5}\Gamma_{\mathbb{A}}(u,t).
\end{aligned}$$

The last inequality follows because we assumed that $\Gamma_{\mathbb{A}}(u,t) < \frac{3}{5}\Gamma(u,t)$. In $(*)$ we used the fact that the activity of all nodes in $N^1(u)$ can only grow from round t to $t+1$, except for those that neighbor heralds or leaders and those that switch to states $\mathbb{M}$ or $\mathbb{E}$. But the former only make up a small percentage of u's

activity mass and the latter do not exist by our assumptions. We therefore in the following assume that $t' \in \{t, t+1\}$ such that $\Gamma_{\mathbb{A}}(u, t') \geq \frac{3}{5} \cdot \Gamma(u, t)$ and u is in state $\mathbb{A}$ in round t'.

To show that in round t', event H_u occurs, we distinguish the two cases given in the lemma statement. We start with the simpler case (a), where in round t, $1 \leq \Gamma(u) < 5\alpha$ and $\gamma(u) = \frac{1}{2}$. The latter implies that $u \notin N(\mathbb{L} \cup \mathbb{H})$, due to the requirements of the lemma that such nodes have low activity. Because no node in $N^1(u)$ switches to states $\mathbb{M}$ or $\mathbb{E}$ in round t and neighbors of leaders and heralds have low activity, u's activity mass cannot change much. Neighbors of leaders and heralds cause a drop of at most $\Gamma(u, t)/100$ and all others increase their activity by at most $\sigma_{\oplus}$. More precisely, $0.99\Gamma(u, t) \leq \Gamma(u, t') \leq (1+\epsilon_{\gamma})\Gamma(u, t)$. We know that $\Gamma_{\mathbb{A}}(u, t') \geq \frac{3}{5}\Gamma(u, t) \geq \frac{3}{5}$. Consequently, since u is in state $\mathbb{A}$, it has activity level $\gamma(u, t') = \frac{1}{2}$, and the total activity mass $\Gamma_{\mathbb{A}}(u, t') - \gamma(u, t')$ of all *neighbors* is between $\frac{1}{10}$ and $\sigma_{\oplus} 5\alpha = O(1)$. Therefore, w.c.p., u and exactly one of its neighbors v operate on channel $\lambda = \mathcal{A}_1$. (Recall that a node w in state $\mathbb{A}$ chooses channel $\mathcal{A}_1$ with probability $\frac{\gamma(w)}{2}$.) Because we assume that u is $\frac{1}{5\alpha}$-fat at time t, $\Gamma(v, t)$ is also bounded and therefore, w.c.p., no other neighbor of v picks channel $\mathcal{A}_1$. Hence, the only thing that is missing to show that event H_u occurs with constant probability is to prove that no other node in $N^2(u)$ hears a message on channel $\mathcal{A}_1$ in round t'.

But this follows from Lemma B.1.(5) by choosing $S = N^2_{\mathbb{A}}(u) \cap (N(u) \cup N(v))$.

We have shown that in case (a), the event H_u occurs with constant probability. Let us therefore switch to case (b), where $N(u, t)$ is η-fat and $\Gamma(u, t) \geq 5\alpha$. For the following argumentation, we define

$$\hat{N}^1_{\mathbb{A}}(u, t') := \left\{ v \in N^1_{\mathbb{A}}(u, t') \setminus N(\mathbb{L} \cup \mathbb{H}) : \Gamma^u_{\mathbb{A}}(v, t') \geq \frac{\Gamma_{\mathbb{A}}(u, t')}{2\alpha} \right\}$$

and

$$\hat{\Gamma}_{\mathbb{A}}(u, t') := \sum_{v \in \hat{N}^1_{\mathbb{A}}(u, t')} \gamma(v, t').$$

To analyze the probability of the event H_u, consider two neighboring nodes $v, w \in N^1_{\mathbb{A}}(u, t') \setminus N(\mathbb{L} \cup \mathbb{H})$. We define $L_{v,w}$ to be the event that in round t' both v and w decide to operate on channel $\lambda := \lceil \log_2 \Gamma(u, t) \rceil$ and no other node in $N(v) \cup N(w)$ chooses the same channel λ. Further $H_{v,w}$ is the event that $L_{v,w}$ occurs and in addition, no node in $N^2(u) \setminus (N(v) \cup N(w))$ receives a message on channel λ in round t'. Lemma B.1.(5) implies again that $\mathbf{P}(H_{v,w} | L_{v,w}) = 1 - O(\pi_\ell)$. Further, note that $H_u = \bigcup_{v,w \in N^1_{\mathbb{A}}(u) \setminus N(\mathbb{L} \cup \mathbb{H}), \{v,w\} \in E} H_{v,w}$, and we have $H_{v,w} = H_{w,v}$ and $H_{v,w} \cap H_{v',w'} = \emptyset$ for $\{v, w\} \neq \{v', w'\}$. It therefore holds that

$$\mathbf{P}(H_u) = \frac{1 - O(\pi_\ell)}{2} \cdot \sum_{\substack{\{v,w\} \in E, \\ (v,w) \in (N^1_{\mathbb{A}}(u,t'))^2}} \mathbf{P}(L_{v,w}). \tag{35}$$

The probability for a node $v \in \mathbb{A}$ to choose channel λ is $\gamma(v) \cdot 2^{-\lceil \log \Gamma(u,t) \rceil} \in \left[\frac{\gamma(v)}{2\Gamma(u,t)}, \frac{\gamma(v)}{\Gamma(u,t)}\right]$. We can therefore bound the probability that $L_{v,w}$ occurs in round t' as

$$\begin{aligned}
\mathbf{P}(L_{v,w}) &\geq \frac{1}{4} \cdot \frac{\gamma(v, t')\gamma(w, t')}{\Gamma(u, t)^2} \cdot \prod_{x \in N(v) \cup N(w)} \left(1 - \frac{\gamma(x, t')}{\Gamma(u, t)}\right) \\
&\geq \frac{\gamma(v, t')\gamma(w, t')}{4\Gamma(u, t)^2} \cdot 4^{-\sum_{x \in N(v) \cup N(w)} \frac{\gamma(x,t')}{\Gamma(u,t)}} \\
&\geq \frac{\gamma(v, t')\gamma(w, t')}{4\Gamma(u, t)^2} \cdot 4^{-\sigma_{\oplus} \frac{\Gamma(v,t) + \Gamma(w,t)}{\Gamma(u,t)}} \\
&\geq \frac{\gamma(v, t')\gamma(w, t')}{\Gamma(u, t)^2} \cdot 4^{-1-\sigma_{\oplus} \frac{2}{\eta} \frac{\Gamma(u,t)}{\Gamma(u,t)}} = \frac{1}{4 \cdot 16^{\frac{\sigma_{\oplus}}{\eta}}} \cdot \frac{\gamma(v, t')\gamma(w, t')}{\Gamma(u, t)^2}.
\end{aligned}$$

The last inequality follows because in round t, node u is η-fat, and therefore

$$\Gamma(u,t) \geq \max\{\Gamma(v,t), \Gamma(w,t)\}.$$

In the following, we restrict our attention to the events $L_{v,w}$ for $v \in \hat{N}^1_{\mathbb{A}}(u,t')$ as these are the only ones for which we obtain a significant lower bound on the probability that they occur. For $v \in \hat{N}^1_{\mathbb{A}}(u,t')$, let $K_v := \bigcup_{w \in N(v) \cap N^1_{\mathbb{A}}(u,t') \setminus N(\mathbb{L} \cup \mathbb{H})} L_{v,w}$ be the event that $L_{v,w}$ occurs for some neighbor w of v. For a node $v \in \hat{N}^1_{\mathbb{A}}(u,t')$, we then have

$$\begin{aligned} \mathbf{P}(K_v) &= \sum_{w \in N(v) \cap N^1_{\mathbb{A}}(u,t') \setminus N(\mathbb{L} \cup \mathbb{H})} \mathbf{P}(L_{v,w}) \\ &\geq \frac{1}{4 \cdot 16^{\frac{\sigma_\oplus}{\eta}}} \cdot \frac{\gamma(v,t')}{\Gamma(u,t)^2} \sum_{w \in N(v) \cap N^1_{\mathbb{A}}(u,t') \setminus N(\mathbb{L} \cup \mathbb{H})} \gamma(w,t') \\ &\geq \frac{1}{4 \cdot 16^{\frac{\sigma_\oplus}{\eta}}} \cdot \frac{\gamma(v,t')\left(\Gamma^u_{\mathbb{A}}(v,t') - \gamma(v,t') - \frac{1}{100}\right)}{\Gamma(u,t)^2} \\ &\geq \frac{1}{4 \cdot 16^{\frac{\sigma_\oplus}{\eta}}} \cdot \frac{\gamma(v,t')\left(\Gamma_{\mathbb{A}}(u,t') - 1.02\alpha\right)}{2\alpha\Gamma(u,t)^2} \qquad (36) \\ &\geq \frac{1}{4 \cdot 16^{\frac{\sigma_\oplus}{\eta}}} \cdot \frac{\gamma(v,t')}{4\alpha\Gamma(u,t)}. \end{aligned}$$

Inequality (36) follows because $\gamma(v,t') \leq \frac{1}{2}$, $v \in \hat{N}^1_{\mathbb{A}}(v)$ and thus $\Gamma^u_{\mathbb{A}}(v,t') \geq \frac{\Gamma_{\mathbb{A}}(u,t')}{2\alpha}$. The last inequality follows from $\Gamma_{\mathbb{A}}(u,t') \geq \frac{3}{5} \cdot \Gamma(u,t)$ and thus $\Gamma_{\mathbb{A}}(u,t') - 1.02\alpha \geq \frac{3}{5} \cdot \Gamma(u,t) - 1.02\alpha \geq \frac{1}{4}\Gamma(u,t)$ if $\Gamma(u,t) \geq 5\alpha$. Using (35), we can now bound the probability of event H_u in round t' as

$$\begin{aligned} \mathbf{P}(H_u) \geq \frac{1 - O(\pi_\ell)}{2} \sum_{v \in \hat{N}^1_{\mathbb{A}}(u,t')} \mathbf{P}(K_v) &\geq \frac{1 - O(\pi_\ell)}{64\alpha \cdot 16^{\frac{\sigma_\oplus}{\eta}}} \sum_{v \in \hat{N}^1_{\mathbb{A}}(u,t')} \frac{\gamma(v,t')}{\Gamma(u,t)} \\ &= \frac{1 - O(\pi_\ell)}{64\alpha \cdot 16^{\frac{\sigma_\oplus}{\eta}}} \cdot \frac{\hat{\Gamma}_{\mathbb{A}}(u,t')}{\Gamma(u,t)}. \end{aligned} \qquad (37)$$

Applying Lemma A.2 to the graph induced by the nodes in $N^1_{\mathbb{A}}(u,t')$, the activity sum of nodes in $\hat{N}^1_{\mathbb{A}}(u,t')$ can be lower bounded as

$$\hat{\Gamma}_{\mathbb{A}}(u,t') \geq \frac{\Gamma_{\mathbb{A}}(u,t')}{2\alpha} \geq \frac{3\Gamma(u,t)}{10\alpha}.$$

Together with (37), this proves that also in case (b), the event H_u occurs with constant probability in a round $t' \in \{t, t+1\}$. Note also that in both cases (a) and (b), for π_ℓ and $\sigma_\oplus$ sufficiently small, the probability that H_u occurs can be lower bounded by a constant C_η that is independent of the probabilities π_ℓ and $\sigma_\oplus$.

To complete the proof, assume that in round t', event H_u occurs with probability C_η and if it occurs, nodes v and w are the two nodes in $N^1(u)$ participating on channel λ (channel $\mathcal{A}_1$ in case (a)). Let M be the event that no herald is created on a channel $\mathcal{A}_i \neq \lambda$ in round t'. Clearly, the probability that M occurs is lower bounded by the probability that no herald is created on any channel in round t'. By Lemma B.1.(1), we therefore have $\mathbf{P}(M) = 1 - O(\pi_\ell)$. For the probability that events H_u and M both occur, we then get

$$\mathbf{P}(M \cap H_u) = 1 - \mathbf{P}(\overline{M} \cup \overline{H}_u) \geq 1 - \mathbf{P}(\overline{M}) - \mathbf{P}(\overline{H}_u) = C_\eta - O(\pi_\ell).$$

Recall that probability C_η is a constant independent of π_ℓ. Conditioned on the event that $M \cap H_u$ occurs, the probability that one of the two nodes v, w listens on channel λ and the other one broadcasts on the channel

is $2\pi_\ell(1-\pi_\ell)$. In that case one of the two nodes becomes a herald candidate and the other one its leader candidate. Also, $M \cap H_u$ implies that in round t' no other herald candidates are created in $N^2(u)$. Let t'' be the round in $\{t, t+1\} \setminus t'$. If in addition in round t'' and in the remaining rounds $t+2, \ldots, t+7$ no herald candidate is created in $N^2(u)$, nodes v and w make it through the handshake and become an isolated leader-herald pair as claimed by the lemma. By Lemma B.1.(1), this happens with probability $1 - O(\pi_\ell)$, which by choosing π_ℓ sufficiently small concludes the proof. $\square$